THE
ECONOMIC
JOURNAL

The Economic Journal, 1–42 https://doi.org/10.1093/ej/ueac042
Advance Access Publication Date: 29 July 2022

ON TARGET? SANCTIONS AND THE ECONOMIC INTERESTS OF ELITE POLICYMAKERS IN IRAN*

Mirko Draca, Jason Garred, Leanne Stickland and Nele Warrinnier

How successful are sanctions at targeting the economic interests of political elites in affected countries? We study the case of Iran, using information on the stock exchange-listed assets of two specific political entities with significant influence over the direction of Iran's nuclear programme. Our identification strategy focuses on the process of negotiations for sanctions removal, examining which interests benefit most from news about diplomatic progress. The results indicate the 'bluntness' of sanctions on Iran, but also provide evidence of their effectiveness in generating substantial economic incentives for elite policymakers to negotiate a deal for sanctions relief.

Economic sanctions are an important tool of foreign policy, providing an instrument by which states may attempt to influence policies abroad without resorting to military force or covert action. In recent years, sanctions have been at the forefront of international responses to Russia's foreign policy decisions regarding the Ukraine and to Iran's programme for the development of nuclear technology. Other prominent recent examples of sanctions have included measures levelled against Burma, Iraq, North Korea and Syria, among many others.

Once imposed, sanctions act as a 'carrot' for policymakers in the sanctioned country, as the actors imposing sanctions offer to remove them in exchange for policy reform. But if imposing sanctions is costly for the senders then they will prefer to design this incentive as efficiently as possible. In theory, sanctions should therefore be focused on the sources of income most valued by those responsible for the key policy decisions, rather than the economy of the receiving country as a whole. In line with this logic, a central principle underlying the design of modern sanctions regimes is that they should be targeted as much as possible to the economic interests of elite decision makers.[1]

* Corresponding author: Mirko Draca, Department of Economics, University of Warwick, Coventry, CV4 7AL, UK. Email: m.draca@warwick.ac.uk

This paper was received on 12 March 2021 and accepted on 13 July 2022. The Editor was Hans-Joachim Voth.

The data and codes for this paper are available on the Journal repository. They were checked for their ability to reproduce the results presented in the paper. The replication package for this paper is available at the following address: https://zenodo.org/record/6638340.

We thank the editor, Hans-Joachim Voth and two anonymous referees, as well as Abel Brodeur and seminar participants at Paris School of Economics, LSE, King's College London, Sussex, Helsinki, Carleton, Ottawa, the NBER Economics of National Security Summer Institute, the EBRD Economics of the Middle East Conference (especially our discussant Simon Commander), the Annual Conference of the Canadian Economics Association, the Royal Economic Society Annual Conference, the Warwick University Peace Studies Association Conference, the LICOS Centre for Institutions and Economic Performance 25 Years Workshop and the Free University of Bozen-Bolzano's Applied Microeconomics Workshop. We also thank Femi Amoo, Terry Cheng, Divya Deepthi and Neel Ocean for excellent research assistance. This research was financially supported by KU Leuven (long-term structural funding—Methusalem funding by the Flemish Government), the Research Foundation Flanders (FWO) (under grant G074718N) and Warwick's Competitive Advantage in the Global Economy (CAGE) research centre.

[1] For example, the European Commission (2008) asserts the principle that '[a]s a general rule, sanctions should target as closely as possible the individuals and entities responsible for the undesirable policies and actions, thus minimising adverse effects on others'. An alternative form of targeting focuses on particular industries, as with arms embargoes (DellaVigna and La Ferrara, 2010).

In practice, however, it may be difficult to design sanctions that stay 'on target'. This might be because the targeted group is particularly well placed to avoid even carefully designed sanctions, or is able to redistribute the effects of sanctions to less powerful groups within the sanctioned country.[2] Alternatively, the policy instruments used by the sender could be too blunt to target any particular group within the sanctioned country. While states are increasingly employing 'smart sanctions' instruments designed to affect only particular firms or individuals, many sanctions regimes also incorporate more comprehensive trade and financial restrictions, which may be difficult to selectively 'switch on and off' across different agents in a receiving country.[3]

In this paper, we consider the case of Iran, which was subject to multilateral sanctions related to its nuclear programme until early 2016. Over several years of diplomatic negotiations ending in July 2015, Iran and its sanctioners came to an agreement in which Iran consented to changes in its nuclear policies in exchange for the lifting of economic sanctions. If these sanctions were successfully targeted, their removal should have benefited the policymaking elite, rewarding them for changing the policy for which sanctions were imposed.

We therefore examine the effects of the lifting of sanctions on Iran, in order to test for the effectiveness of targeting in terms of the incentives faced by Iranian policymakers. We consider two key actors who were targeted by the senders of sanctions because of their important roles in decisions about Iran's nuclear programme: the Islamic Revolutionary Guard Corps (IRGC) and Iran's Supreme Leader, Ali Khamenei. Both of these actors reportedly control large conglomerates, and we check whether certain observable assets of those conglomerates—specifically, their holdings in firms listed on the Tehran Stock Exchange (TSE)—were positively affected by diplomatic progress towards a multilateral sanctions deal. We also compare the TSE returns of these 'target firms' to those of a set of unrelated firms with no known connection to the targeted conglomerates, so as to better understand the relative impact of sanctions relief across these two (TSE-listed) parts of the Iranian economy.

To accomplish this, we identify 'information shocks' regarding the progress of diplomatic negotiations, and compare contemporaneous shifts in firm-level TSE returns across target and unrelated firms. This research design based on sanctions *removal*, in addition to being motivated by theory, has key advantages for identification. First, it allows us to estimate the impact of the full suite of Iran sanctions rather than individual measures, because news about diplomatic progress relates to the lifting of the entire sanctions regime. In contrast, designs focusing on the introduction of individual sanctions measures are subject to the heterogeneity of these measures, both in terms of the exact targets (for example, different entities are targeted by different policy instruments) and the size of the specific effects (some measures may be more effective than others). Second, it allows us to exploit unexpected and precisely timed progress in diplomatic negotiations.

We first explore a compelling case study: the stock market reaction to a breakthrough in multilateral negotiations in Geneva between Iran and the main sanctioning countries in November 2013. This 'big bang' event is conveniently timed, with much information about the progress of negotiations accruing over the TSE's weekend break. Because of this, we are able to cleanly

[2] For example, the UN's oil-for-food programme in Iraq was designed to maintain an oil embargo on Iraq while still providing the country with access to humanitarian supplies, but there is strong evidence that the policy was circumvented, presumably to the benefit of Iraq's political elite (Hsieh and Moretti, 2006).

[3] Such a problem of 'bluntness' leading to unintended consequences is not unique to sanctions policy and may also be present in strategies involving armed intervention (Dell and Querubin, 2018) or military assistance (Dube and Naidu, 2015).

capture the TSE market response to the most important turning point in the negotiations for sanctions removal.

We then use high-frequency text-based measures of news about sanctions negotiations to capture information shocks covering the full period from the re-opening of serious negotiations in 2012 to agreement on a final deal in 2015. This allows us to test the sensitivity of different types of firms to news about possible progress towards the lifting of sanctions using a large collection of events. The measures we use—derived from the Factiva and GDELT (Global Database of Events, Language and Tone) databases—provide quantitative information on coverage based on an extensive library of news sources.[4]

Our results show a consistent pattern of responses by the two sets of firms to these information shocks. Specifically, in each of our empirical exercises, we find that the stock returns of firms owned by targeted political groups (our target portfolio) and companies unrelated to these groups (our non-target portfolio) both react positively to information, indicating progress in diplomatic negotiations. However, the effects are significantly larger for the target portfolio. This pattern is compatible with a scenario in which sanctions rely on 'blunt instruments' and so affect the economy of the receiving country more widely than intended. But it also suggests that the economic interests of Iran's political elite—of which we observe a set of listed assets constituting 10% of the TSE's total market capitalisation—were indeed affected by sanctions, and differentially so when compared with other listed firms.

Moreover, we find that these impacts were economically significant. The November 2013 breakthrough in Geneva resulted in a rise of 3.8 percentage points in stock returns for elite-owned firms over two days, and so is estimated to have increased the value of IRGC and Setad assets by approximately $465 million in total. Furthermore, both the 'Geneva effect' and the average impacts of multiple discrete shocks we identify from the news coverage data are of a similar magnitude to the upper tail of average firm returns. These shocks are also large as compared to the shifts in returns induced by movements in oil prices.

We also show that a later negative shock in expectations of the sanctions deal's continuation, provided by Donald Trump's surprise election victory, is similarly associated with large differential stock returns (negative in this case) for target firms. Subsequent developments, again captured using daily news coverage measures, show no additional sensitivity for target portfolio returns, possibly implying that sanctions policies went 'off target' after Trump took power. However, the events captured by our coverage measures during the period after the US election (such as new sanctions in response to Iranian missile tests) are quite different from the diplomatic negotiations underlying our main results.

Finally, we show evidence that our main results are not driven by a potential alternative explanation: that diplomatic progress also changes the likelihood of conflict between Iran and other countries, which could itself have heterogeneous effects on target and non-target firms. We first observe that, while both portfolios are sensitive to a political betting market measure of Iran's probability of direct military conflict with the United States or Israel, the target portfolio is not differentially responsive to this information. We then find that the returns of arms industry firms listed on stock markets outside Iran show no significant co-movement with news of diplomatic progress towards a deal on Iran sanctions.

[4] This aspect of our study builds on recent contributions that have incorporated text-based information on salient news shocks into stock market studies (e.g., Tetlock, 2007; Loughran and McDonald, 2011; Baker *et al.*, 2016).

Overall, our finding of systematic co-movement between target firm returns and events relating to sanctions removal provides evidence that multilateral sanctions against Iran succeeded in inducing some degree of 'income targeting' of the political elite. Even though we cannot say whether decision makers were explicitly driven by this incentive when negotiating the removal of sanctions, our empirical evidence suggests that such an incentive existed.

Related literature. Our study makes a novel contribution to the substantial literature on the economics of sanctions. A large body of empirical work has studied the question of whether, and under what circumstances, sanctions accomplish the stated goals of the sender. Much of this work builds on the cross-country analysis and dataset of Hufbauer *et al.* (1990) and focuses on whether the offending policies of sanctioned countries are actually altered after sanctions are imposed. Very few papers zoom in to consider whether the incidence of sanctions within a given country is in line with the intentions of the sender.

Recent firm-level research by Ahn and Ludema (2019; 2020) finds that US and EU 'smart sanctions' on specific Russian companies, and their shareholders, managers and directors, have had substantial negative effects on those firms' performance. Their identification strategy is based on the timing of the introduction of sanctions and uses outcomes from annual firm accounts. Our paper chiefly differs in its focus on the economic interests of elite policymakers as an incentive central to the operation of modern sanctions regimes.[5] Also, our research design based on the lifting of sanctions allows us to assess impacts across the full set of sanctions imposed on Iran, including both 'smart' and comprehensive measures. Finally, we take advantage of daily stock price data in order to use high-frequency variation for identification.

Within the wider political economy literature, our paper is one of a number of recent studies exploiting variation in stock market returns within an event study framework. This methodology has been applied to topics such as the private benefits of civil conflict (Guidolin and La Ferrara, 2007), the economic implications of covert foreign intervention (Dube *et al.*, 2011) and the value of political connections (Fisman, 2001; Faccio, 2006; Knight, 2006; Coulomb and Sangnier, 2014; Acemoglu *et al.*, 2018). We also contribute to a growing literature on the political economy of the Middle East (e.g., Jaeger and Paserman, 2008; Berman *et al.*, 2011, Durante and Zhuravskaya, 2018).

Structure. The paper is organised in the typical way. Section 1 presents a conceptual framework to motivate our study and its research design. Section 2 provides background information on the history of sanctions on Iran and the political entities targeted by these sanctions. Section 3 describes the data we use, and Section 4 then presents our empirical analysis. Section 5 assesses the possible alternative interpretation of our results in terms of changes in the risk of conflict. Section 6 offers concluding remarks.

1. Conceptual Framework

In this section, we first set out a simple static framework to clarify why countries might want to target the economic interests of elite policymakers, and why the effectiveness of such targeting is an important economic question. As part of this, we frame the effects of sanctions as creating an

[5] Some studies have considered the impact of sanctions on political outcomes (rather than regime assets): Marinov (2005) found that leaders of countries subject to economic sanctions are more likely to lose power, while Allen (2008) noted a positive relationship between sanctions and anti-government activity. Another paper considering the economic impact of sanctions on subgroups within a country is that of Neuenkirch and Neumeier (2016), who argued that US sanctions have led to an increase in poverty in sanctioned countries. Haidar (2017) found heterogeneous effects of sanctions on Iran on exporting firms depending on characteristics such as exporter size.

incentive ('carrot') for policy changes in a sanctions-receiving country, distinguishing between targeted and non-targeted subgroups within the receiver. We then discuss the effect of 'blunt instruments' on the realised incidence of sanctions and show that the cost of sanctions tends to increase under this scenario, due to the wider distribution of impacts across these subgroups. Finally, we consider the implications of time-varying progress in negotiations towards sanctions relief, in order to motivate our empirical strategy and aid in the interpretation of our results.

For simplicity, we assume in this framework that the costs imposed by the sanctions to the receiver at the point of introduction are symmetric with the benefits of removing the sanctions. This necessarily abstracts from more complex scenarios where the size and incidence of costs that arise when sanctions first hit are different from the benefits when sanctions are lifted. In practice, the main focus of our study is on the latter, since our goal is to investigate the economic incentives for elite policymakers to negotiate sanctions relief.[6]

1.1. *Set-Up*

We consider a scenario in which sanctions are imposed in order to induce changes in a particular policy in the sanctioned country, since this was the usual publicly stated goal of multilateral sanctions on Iran.[7] Consider two countries, a sender S of sanctions and a receiver R. Within R, a group p (the 'political elite') has control over a policy that, if in place, benefits p but imposes a cost on S. In response, S imposes economic sanctions on R, offering to lift these sanctions if R ends the policy. While sanctions may also be costly to the sender S (because of lost trade, transactions costs or political considerations), we assume for simplicity that these costs are lower than the cost to S of the offending policy in place in receiver R.

Once sanctions are imposed, they work as a carrot rather than a stick: as long as group p within R benefits from sanctions relief, the removal of sanctions is effectively a reward for ending the policy that harms S. Say that Y_R represents the aggregate income of R, an amount Y_p of which accrues to the political elite p, with all others in R (group o) receiving the remainder, Y_o. Also assume that group p has a welfare function consisting of three additive terms: b_p, the benefit to p from the policy being in place; national income Y_R and an additional weight $\alpha - 1$ on the group's own income Y_p. Then welfare of p is equal to $b_p + \alpha Y_p + Y_o$. The relevance of Y_o to the welfare of p might be due to altruism or 'good governance' by group p, or because of the need for support from the population as a whole in order to sustain the political regime. But, as long as $\alpha > 1$, group p values increase in its own income more highly than similar increases in the income of the rest of the population. Finally, say that sanctions incur a cost to R that manifests as lost income y, and that this cost is spread additively across group p and group o, so that $y = y_p + y_o$.

To understand the potential attractiveness of targeting to S, assume that the cost of sanctions to S also rises in y, so that increasing the cost of sanctions in the receiving country also increases the cost to the sender.[8] Then consider the scenario in which S has complete control over their

[6] Note that we also abstract from the stage in which sanctions are threatened but not yet imposed; see Eaton and Engers (1992) for a model that incorporates this stage.

[7] The framework in this section can easily be extended to accommodate other scenarios; for example, to show that the targeting of sanctions may also be desirable if the sanctioners' goal is to alter the balance of power between different actors within the sanctioned country, as in Kaempfer and Lowenberg (1988).

[8] A sanctions regime of higher intensity might, for example, lead to greater foregone gains from trade or increased transaction costs for the sender.

incidence across o and p. In this case, S solves the problem

$$\min_{y_p, y_o} \quad y_p + y_o \quad \text{subject to} \quad \alpha y_p + y_o \geq b_p, \quad y_p \geq 0, \quad y_o \geq 0.$$

The sender must assure that the benefit to group p if the sanctions are lifted is at least as large as the cost of ending the policy, while minimising the cost of sanctions to S itself. Given that $\alpha > 1$, the solution to this problem is to set $y_p = (1/\alpha)b_p$, $y_o = 0$. In other words, as long as each additional unit of income lost by p has a greater effect on the welfare of p than a unit of income lost by others in R, the minimum-cost sanctions policy is to target sanctions so that they affect the income of p exclusively.

1.2. *'Blunt Instruments'*

A fully targeted sanctions regime might not be available to the sender in practice. One possibility is that the sender might have to resort to instruments with some direct effects on group o. This could be due to the unavailability of precisely targeted instruments whose direct effects accrue only to group p, or the fact that the full suite of precisely targeted instruments is insufficiently strong to push policymakers in the receiver to change the policy of interest. Alternatively, even if the sender is able to construct a regime such that the direct effects of sanctions entirely accrue to group p, there might be spillovers within the economy of the receiver. These might include input-output relationships between the groups, or might be due to the ability of group p to shift some of the costs of sanctions to group o. We can summarise all of these possibilities as the 'blunt instruments' problem.

Say that due to blunt instruments, the 'maximally targeted' sanctions regime of cost y actually leads to incidence $y_p = \beta y$, $y_o = (1 - \beta)y$ for some $\beta \in [0, 1]$. The β parameter can therefore be seen as the share of income costs borne by the political elite group p within the receiver country. In simple terms, $\beta = 0$ corresponds to the elite group p escaping all of the income costs of sanctions, while $\beta = 1$ implies that the full income costs fall on the political elite only.

In the 'bluntness' case where $\beta < 1$, the chosen sanctions regime will be costlier than in the scenario where S can control sanctions incidence. In particular, the solution for S is now to set $y = (1/(\alpha\beta + (1 - \beta)))b_p$, which is a greater total cost than under the targeted regime, where $y = (1/\alpha)b_p$. In this model, both S and R (in aggregate) therefore benefit from the ability of S to target sanctions. The effectiveness of targeting—here quantified as the magnitude of β—is thus an interesting economic question.

1.3. *Negotiations for Sanctions Removal*

Once a given sanctions regime is imposed, the above static framework requires that a cost-benefit analysis by group p should lead to an immediate decision to remove or continue the policy. In practice, however, a sender and receiver of sanctions can engage in a period of negotiations over the actual concessions to be made by the receiver, and the sender might escalate or ease the sanctions regime during this period. Based on the progress of negotiations, the expected discounted costs of sanctions and benefits from the offending policy may therefore evolve over time, due to changes to the current situation as well as expected future events. For example, if S and R take a step towards a deal to exchange policy reform for sanctions relief, this will lead to a decrease in the expected (discounted) future costs of sanctions to both p and o and the expected future benefit of the policy to p. Therefore, if we are interested in learning about

β using information from progress in political negotiations, we need an empirical strategy that separates the effects of these negotiations on the costs of sanctions to the targeted political elite p and to the non-elite o.

In this study, we use stock return data to examine the evolution of the asset values of groups corresponding to p and o during negotiations between Iran and its sanctioners. If the values of assets of p co-move positively with news of progress in negotiations, and these asset values are unrelated to any benefits from the policy itself, then we interpret this co-movement as evidence that $\beta \neq 0$. Practically, this implies that the group p benefits from expected sanctions relief because it has not escaped bearing an income cost from the sanctions.

Though our main focus below will be on the assets of elite policymakers, we also provide a broad empirical assessment of the bluntness of sanctions across the publicly listed firms we study. Specifically, we derive further information about β within this context by comparing the relative magnitudes of changes in the values of the two groups' assets. If the values of assets of the non-elite group o are moved by progress towards sanctions relief, we consider this to be evidence that $\beta \neq 1$; that is, the elite group does not bear the full costs and there is some sharing of costs between p and o. We further discuss the mapping between β and our empirical difference-in-difference model in Subsection 4.1.

2. Background

2.1. *Sanctions on Iran and Political Negotiations*

While the United States has maintained economic sanctions on Iran since soon after its 1979 revolution, robust multilateral sanctions prompted by Iran's nuclear programme were imposed only from the mid-2000s.[9] The United Nations Security Council first passed a resolution threatening Iran with sanctions in July 2006, in reference to International Atomic Energy Agency (IAEA) reports stating that the IAEA was unable to determine that Iran's nuclear programme had no military dimension. The resolution called for Iran to suspend enrichment and reprocessing activities associated with its nuclear programme. Sanctions were then imposed in December 2006 and tightened in two subsequent resolutions in 2007 and 2008, with these resolutions specifying that the sanctions would be removed once Iran met requirements set by the IAEA and the Security Council itself.

The sanctions put in place during this period constituted both narrowly focused 'smart sanctions' and more comprehensive measures. For instance, the Security Council resolutions asked states to freeze the assets of specific firms, individuals and other organisations directly involved in Iran's nuclear programme, and also to 'exercise vigilance' over the foreign activities of Iranian financial institutions in general. Sanctions were publicly portrayed as targeting the economic interests of actors with political influence inside Iran, by cutting Iran off on margins that would specifically affect those interests.[10] However, some observers argued that sanctions were ineffective, or even strengthened defenders of the nuclear programme within Iran: for example, the

[9] US sanctions on Iran have included a variety of measures relating not only to weapons proliferation, but also to other issues such as human rights. Our study focuses only on sanctions related to Iran's nuclear programme.

[10] For example, after the United States ended access to the American financial system for a major Iranian bank, the US Treasury's Undersecretary for Terrorism and Financial Intelligence argued that '[w]hile those who are currently benefitting from Iranian integration into the global economy are the ones who will feel this isolation the most, they are also in the best position to persuade the regime that its current track will undermine the future of the Iranian people' (Higgins, 2006).

Wall Street Journal suggested in 2006 that US pressure on Iranian banks ‘ended up boosting the very hard-line forces there that the US wants to curb’ (Higgins, 2006).[11]

Several rounds of international negotiations were held alongside this gradual tightening of sanctions in the wake of the initial 2006 Security Council resolution. The talks with Iran were led by the ‘P5+1’ group of countries: the five permanent members of the UN Security Council (China, France, Russia, the UK and the United States) as well as Germany. This process failed to reach agreement on sanctions relief in exchange for policy changes in Iran.

From 2010 through early 2012, the sanctions regime was greatly intensified, both due to a new UN Security Council resolution in mid-2010 and through various additional measures taken by individual actors including the United States and EU. These new sanctions imposed severe limitations on Iran’s international financial access; for example, the Belgium-based organisation in charge of international banking transactions (SWIFT) removed a number of Iranian banks from its system in early 2012, after pressure from the EU. At the same time, new restrictions were placed on imports of various goods from Iran, including oil, by the EU and others. Meanwhile, the pace of diplomatic negotiations slowed; the only round of high-level multilateral talks during this period, in late 2010 and early 2011, did not produce a breakthrough.

After a hiatus of more than a year, the P5+1 negotiations on Iran’s nuclear programme resumed with a meeting in Istanbul in April 2012, and were characterised as successful by both sides (BBC, 2012). After a series of meetings over the following months, the first major diplomatic breakthrough occurred in Geneva in November 2013, when the parties agreed to a framework agreement. This deal was followed by lengthy negotiations on a final agreement to lift sanctions in exchange for concessions related to Iran’s nuclear programme. The interim agreement reached at Geneva, which was originally due to expire in July 2014, was extended twice, with an eventual final deadline of June 2015. The framework of a final agreement was reached during high-level negotiations in Lausanne in April 2015, and the details of this deal were eventually concluded on July 14, 2015.

Importantly for our identification strategy, diplomatic progress towards a final agreement from 2012 to 2015 was uneven and fraught with uncertainty. Media reports from various points of the timeframe we study document unexpected breakthroughs (e.g., Warrick and Rezaian, 2013a) as well as long periods of slow progress (e.g., Pawlak and Dahl, 2014) and doubts in advance of negotiated deadlines (e.g., Borger, 2015). Moreover, the eventual success of the process was far from certain in the early part of this period, when there were reports of ‘long odds’ of ever reaching a deal (Warrick and Rezaian, 2013b).

After the agreement was officially approved by all parties and the IAEA reported that Iran had met its commitments under the deal, multilateral sanctions were lifted in January 2016. However, the situation changed when in November 2016, an opponent of the agreement, Donald Trump, won the US presidential election. While the other members of the P5+1 group maintained their support for the accord, the United States withdrew from the deal and began reimposing sanctions in May 2018. US nuclear-related sanctions on Iran were fully restored in November 2018. These new US sanctions were not accompanied by high-level political negotiations with Iran over conditions for their removal.

[11] The article suggests that financial sanctions negatively affected firms unconnected to the regime, creating opportunities for hardliners to purchase such companies; this could be interpreted as an especially problematic case of the ‘blunt instruments’ issue we have outlined.

2.2. *Targeted Political Entities*

Throughout the development of multilateral sanctions on Iran, the UN, EU and United States have publicly linked sanctions policies to specific entities who have a major political influence over the direction of Iran's nuclear strategy, but also hold substantial economic assets. We discuss the nature of each of these actors' influence over the nuclear programme and the structure of their economic interests in turn.

2.2.1. *Islamic Revolutionary Guard Corps*

The IRGC is a military organisation with the explicit political role of guarding the 1979 Islamic revolution and promoting it outside Iran. The IRGC therefore officially functions as a branch of Iran's armed forces alongside its regular military, but is also known to be influential in politics, and former IRGC members such as Mahmoud Ahmedinejad (president of Iran between 2005 and 2013) frequently serve in prominent political roles.

The United States, EU and the UN Security Council have all claimed the existence of close links between the IRGC and Iran's nuclear programme in official documents. For example, in the 2010 resolution mentioned earlier, the UN Security Council noted 'with serious concern the role of elements of the [IRGC] ... in Iran's proliferation sensitive nuclear activities and the development of nuclear weapon delivery systems' (UN, 2010). In a document from the same year, the EU characterised the IRGC as 'responsible for Iran's nuclear programme' (EU, 2010).

As noted above, the multilateral sanctions regime included instruments intended to influence political actors through comprehensive restrictions (e.g., on financial transactions), as well as smart sanctions incident on specific entities. The IRGC's extensive economic presence in Iran—originating from its participation in military industries and provision of services for current personnel and veterans, but now covering a wide range of industries—makes it potentially vulnerable to sanctions of both types. Indeed, beginning with the first UN Security Council resolution imposing sanctions in 2006, individuals and firms associated with the IRGC have been the subject of a large number of smart sanctions. For instance, the Fordow fuel enrichment plant became a focus of nuclear diplomacy after its disclosure in September 2009, and the UN Security Council's 2010 resolution on Iran included the imposition of smart sanctions on a group of IRGC-owned firms that had been involved in its construction.

2.2.2. *Supreme Leader and Setad*

While Iran's political system includes an elected president as well as a parliament, it is the Supreme Leader, Ali Khamenei, who is effectively its most powerful figure. This includes ultimate political control of its nuclear programme, as recognised by Wendy Sherman, the lead negotiator for the United States in the Iran talks, in a Congressional hearing in December 2013: 'The Supreme Leader is the only one who really holds the nuclear file, [and] makes the final decisions about whether Iran will reach a comprehensive agreement to forgo much of what it has created in return for the economic relief it seeks' (Kerr, 2018).

As with the IRGC, the Supreme Leader's economic interests have been recognised and targeted by the senders of sanctions on Iran. The EU imposed sanctions on Mohammad Mokhber in 2010, characterising him as 'President of the Setad Ejraie Foundation, an investment fund linked to Ali Khamenei, the Supreme Leader' (EU, 2010). In June 2013, the United States Treasury announced sanctions on 'a major network of front companies controlled by Iran's leadership' (US Treasury, 2013). The Treasury identified thirty-seven firms in a network associated with the organisation

EIKO (the Execution of Imam Khomeini's Order), also known as Setad, which it claimed were 'generating billions of dollars in profits for the Iranian regime each year'. While the US Treasury did not provide further details on Setad's political ties, a Reuters investigative report on Setad later in 2013 identified it as a conglomerate controlled by Supreme Leader Khamenei, and noted that Setad's revenues allowed for Khamenei's financial independence from Iran's parliament and national budget process (Stecklow *et al.*, 2013).

A report in 2015 revealed that the targeting of the Supreme Leader's economic base in Setad was an explicit political strategy on the part of sanctioners. Officials quoted by Reuters (Torbati and Dehghanpisheh, 2015) indicated that the organisation was hit with smart sanctions in 2013 because the United States 'saw it as close to Khamenei and believed that the sanctions might induce him back to serious nuclear negotiations'. Another official also noted that at that time, nuclear talks were deadlocked and action was felt to be needed: 'The reason why we dropped the hammer on them when we did is because we were attempting to put pressure on the Supreme Leader (to agree to a deal)' (Torbati and Dehghanpisheh, 2015). This report also suggested that the lifting of sanctions yielded tangible economic benefits for the conglomerate, such as the facilitation of international business deals. In Subsection 4.2, we take a closer look at the potential relevance to Setad of the 2013 US smart sanctions relative to other (more comprehensive) sanctions instruments that were already in place.

3. Data

3.1. *Stock Return Data*

We collect information on stock returns of listed Iranian firms by web scraping daily data (TSE, 2018) from the website of the Tehran Stock Exchange (TSE). As background, it should be noted that, while the TSE is not as deep and sophisticated as many North American and Western European stock exchanges, it provides a suitable setting for conducting an event study. According to the World Bank's Global Financial Development database (see Čihák *et al.*, 2012), the TSE puts Iran in the second (from top) quartile of financial access (share of companies outside the top ten in market capitalisation), the third quartile of financial depth (market capitalisation as a share of GDP) and the second quartile of financial efficiency (shares traded as a proportion of total market capitalisation), relative to other countries.[12] The TSE trades for three hours per day (9:00 a.m. to noon), five days per week (Saturday through Wednesday).

Our main sample period covers the revived multilateral negotiations leading up to the final deal. The sample therefore begins with the first day of negotiations in Istanbul (which happened to fall on the first day of the TSE's trading week) on Saturday, April 14, 2012, and ends on Wednesday, July 15, 2015 (the last day of the TSE's trading week), the day after a final agreement was reached in Vienna. In brief, we choose this window because April 2012 represents the beginning of an uninterrupted phase of negotiations (that is, a period when there was no major breakdown or hiatus in talks) and because the majority of sanctions were in place by early 2012. As we outline in Subsection 4.2, this latter point facilitates an event study research design centred on the effects of fully lifting sanctions.

[12] The TSE ranks very similarly to Egypt—the context studied by Acemoglu *et al.* (2018)—by these measures, based on data from the Egyptian Exchange. A study of market efficiency (Jahan-Parvar and Mohammadi, 2013) indicates that the TSE is CAPM efficient at the monthly level and displays patterns of international integration comparable to Middle Eastern markets that are considered to be open, such as Bahrain, Israel and Turkey.

As of the beginning of our sample period, there were 325 firms listed on the TSE. As discussed in the next subsection, our baseline sample is made up of two firm portfolios, which together include 138 of these 325 firms. Daily returns for each firm are calculated by subtracting a stock's closing price for the previous day on which the stock was traded from its closing price for the current day and dividing by the closing price for the previous day. We multiply daily returns by a hundred so that they are expressed as percentages. We also exclude the top and bottom 1% of all observed returns from our sample, so that our results are not driven by outliers due to measurement error or very large positive or negative returns.[13]

For our exercises relating to the 2016 US election and subsequent events, we instead work with a sample period beginning sixty trading days before the US election and ending on the day when news of the US withdrawal from the deal reached the TSE. This period covers August 13, 2016 to May 9, 2018. As with our main sample period, we drop the top and bottom 1% of observed returns over this time.

3.2. *Ownership Data*

Our goal is to assemble a portfolio of firms listed on the TSE that were assets of the IRGC and/or Setad during the sample period in order to gather evidence on whether the political actors controlling these conglomerates may have gained from evolving progress towards sanctions relief. In our preferred approach, we use information on firms explicitly identified as IRGC and Setad assets by the main 'senders' of sanctions (the United States, UN and EU). We then identify the TSE-listed assets of these members of the IRGC and Setad conglomerates.

As we discuss below, another feature of our strategy is that we identify not only the set of assets on the TSE that are most closely linked to the key targeted conglomerates, but also the TSE firms that are most removed from these conglomerates with no known affiliation via ownership. This is useful for our research design since this set of 'most removed' firms are likely those that the senders of sanctions were least interested in targeting. Hence, they provide the cleanest possible comparison group for an assessment of the bluntness of sanctions across publicly listed firms.

We now proceed to provide details of our baseline procedure for classifying firms, a graphical summary of which is displayed in Online Appendix Figure A1. In Section 4, we also suggest several alternative approaches to the construction of our firm portfolios, and present results of robustness checks using these alternative portfolio definitions.

3.2.1. *Target portfolio*

The initial 'universe' of firms that we consider are those that were listed on the TSE at the start of our main sample period on April 14, 2012, representing a total of 325 firms. As a first step, we exclude firms with business operations directly relating to Iran's nuclear programme, as identified by the senders of sanctions, with prominent examples being Iran's national maritime carrier, IRISL Shipping Lines, and several large banks. These firms are dropped in order to avoid

[13] Our main results are similar if we instead take the less conservative approach of dropping the top and bottom 0.1% of returns, or alternatively if the data are winsorised. In practice, especially large calculated returns tend to be preceded by periods where we do not observe trading in a given stock, implying that these cases do not represent actual day-to-day returns. Our main results also remain similar when we only drop firm days preceded by at least one trading day (or alternatively at least two trading days or one week) where we do not observe that firm's stock being traded. The main results are also robust to dropping the bottom 25% of thinly traded stocks, and to controlling for the number of TSE-wide days off before a trading day.

confounding the expected effects of sanctions relief with anticipated changes in the business of supplying goods or services for the nuclear programme.[14]

In the second step, we begin defining the target group of firms. This is challenging because to do so, we require reliable information on the assets of the IRGC and Setad conglomerates, which is not available on a systematic basis in documents such as annual reports. However, as discussed in Subsection 2.2, Iran's sanctioners have often imposed smart sanctions on individual entities linked with the IRGC or Iran's Supreme Leader. In each case, the rationale for the imposition of sanctions has been laid out in an official document produced by the sanctioner. We can therefore use these documents to identify entities that the UN, EU or United States have explicitly stated are owned or controlled by the IRGC or Setad.

Entities on which US smart sanctions were imposed due to their links with the IRGC may be found using the Department of the Treasury's Specially Designated Nationals and Blocked Persons list (SDN list), in which these organisations are assigned a specific 'IRGC' identifier. We draft a list of all entities that have been tagged with this identifier and subsequently cross-check these with US Treasury press releases that provide further details on the reasons for their inclusion in the SDN list. We add entities into our target group only if the relevant press release states that the entity is owned or controlled by the IRGC.[15] On the other hand, entities sanctioned by the United States due to their links with Setad do not have a separate identifier on the SDN list. However, US smart sanctions on Setad assets were imposed in a single episode in June 2013, and so we use a detailed US Treasury press release from this episode to identify these firms.

Similarly, European Council decision 2010/413/CFSP and its amendments provide lists of EU-sanctioned entities, while UN Security Council resolutions identify organisations designated for smart sanctions by the UN. Both sets of documents provide rationales for these decisions. Again, we only include entities that these records state are owned or controlled by the IRGC or Setad.[16]

The third step in defining the target portfolio involves identifying all TSE-listed assets of the IRGC and Setad entities defined in the process just outlined. Both the IRGC and Setad conglomerates are comprised of a mixture of stock market listed and private, unlisted companies. Thus, although information from the United States, EU and UN allows us to identify seventy-five firms and other entities partly or fully owned by the IRGC and/or Setad, only six of these are themselves listed on the TSE as of the beginning of the sample. These six companies represent approximately 5.5% of the total market capitalisation of the TSE, or $6.3 billion.

To deepen our coverage of IRGC and Setad assets, we also use shareholder information as of the beginning of the sample period, extracted from the TSE website (TSE, 2017), to identify firms in which at least one of the seventy-five targeted entities is a shareholder. In this step, we identify an additional forty-four TSE-listed firms in which the IRGC and/or Setad have ownership stakes, resulting in a target portfolio of fifty firms in total. This set of forty-four

[14] We define this set of firms using a similar approach to the one we use to identify IRGC and Setad assets, discussed in the next paragraph. Specifically, we exclude any firm on which smart sanctions were imposed by the United States, EU and UN for involvement in supplying the nuclear programme, according to official documents. In terms of the conceptual framework, we do this so as to isolate changes in asset values that are unrelated to any benefits from the policy itself; see Subsection 1.3.

[15] Sanctions have been imposed on some entities for providing support or services to the IRGC. For example, Iran's national air carrier, Iran Air, was sanctioned for the transportation of military-related equipment on behalf of the IRGC.

[16] In 2010, the European Union also sanctioned Sina Bank, with the rationale that it was 'very close to the interests of the Daftar (Office of the Supreme Leader). It contributes in this way to funding the regime's strategic interests' (EU, 2010). Because Sina Bank was not described as a Setad asset, we have not included it in the target portfolio. However, the main results are robust to adding Sina Bank and its TSE-listed assets as target firms.

companies makes up about 16.1% of the TSE's market capitalisation, a total of $18.5 billion. The observed share of ownership of these firms by the IRGC and Setad represents about a third of this amount, approximately $5.9 billion. Including the six firms discussed above, we therefore link listed assets constituting 10% of the TSE's total market capitalisation to the IRGC and Setad conglomerates.

3.2.2. *Non-target portfolio*

We next construct a portfolio composed of the listed firms that we can most confidently assume are not assets of either the IRGC or Setad. Our starting point is the 260 listed firms that were not identified by the classification procedure documented in the previous subsection. However, while this excludes the fifty firms with the most direct known ownership connection to the IRGC or Setad, there are two additional groups of firms with more tenuous or less certain connections to these two conglomerates, and we drop each of these from the non-target portfolio.

The first group of companies that we drop are those with what could be termed 'second-degree' connections to the IRGC or Setad. These are the 144 listed firms that are two or more ownership layers below the seventy-five entities identified as IRGC or Setad assets by the United States, EU and UN. We then eliminate a second group of firms that sources other than the United States, EU or UN suggest might be assets of the IRGC or Setad. These sources are a 2013 Reuters investigative report on Setad (Stecklow *et al.*, 2013) and a 2010 American Enterprise Institute report on the IRGC's involvement in the Iranian economy (Alfoneh, 2010).[17] This eliminates another twenty-eight firms from the initial pool, leaving eighty-eight firms in our non-target portfolio.

Table 1 shows summary statistics for firms in the target and non-target portfolios. According to data from the TSE, these two groups of firms make up a similar share of total TSE market capitalisation: between one-fifth and one-quarter in each case. The mean and SD of the daily return variable by firm day are also very similar across the two portfolios. Accounting data from Orbis for 2012 (Bureau van Dijk, 2020), which are available for only a subset of these firms, suggest that there are no statistically significant differences between the mean turnover, assets and labour force of the firms in each portfolio. The industrial composition of the two groups is somewhat different: for instance, a larger proportion of target firms operate in the financial sector.[18] We account for these differences in sectoral composition in some of our empirical specifications below.

3.3. *News Coverage and Event Data*

As part of our research design, we construct daily measures of news coverage of sanctions-related negotiations involving Iran. The main aim of these coverage variables is to provide consistent measures of sanctions-related events that may be salient to TSE market actors during this period of diplomatic progress towards a sanctions deal.

[17] We begin by dropping listed firms that are identified in these sources but are not in our target portfolio. We then use the TSE shareholder data again to construct layers of ownership below each of the entities named in these additional sources, and remove all firms within this ownership structure from the non-target portfolio.

[18] We use a listing of firms by sector in TSE (2011) to classify firms into industries. This provides thirty-seven sectors, mostly corresponding to industries in the NACE classification (Eurostat, 2008) at various levels of aggregation. Many of these sectors are sparsely populated, so to allow for within-industry comparisons of firms, we amalgamate sectors that are adjacent in the NACE classification until there is at least one target and one non-target firm in each industry. This procedure results in fifteen industries (see Table 1), as well as two additional sectors with no target or non-target firms.

Table 1. *Summary Statistics—Sample Firms.*

	Target firms	Non-target firms
Number of firms	50	88
Daily return	0.134	0.136
	(2.01)	(2.05)
Market capitalisation	12,357.47	7,908.29
	(31,273.68)	(22,461.24)
Share of TSE market capitalisation	21.53	24.26
Share of firms by industry:		
Mining	2.0	4.6
Food products	4.0	6.8
Wood/paper/textiles	2.0	3.4
Refined petroleum	4.0	1.1
Chemicals	2.0	10.2
Pharmaceuticals	6.0	10.2
Rubber/plastic/mineral products	28.0	18.2
Basic metals	10.0	10.2
Metal products	2.0	1.1
Electronics/electrical equipment	4.0	8.0
Machinery	4.0	2.3
Motor vehicles	12.0	8.0
Transportation/telecom	2.0	1.1
Finance	16.0	8.0
Construction/real estate	2.0	6.8
Orbis data:		
Number of firms	32	57
Turnover	3,566.48	5,034.33
	(8,899.41)	(23,743.09)
Assets	5,674.74	3,482.93
	(16,403.98)	(9,583.67)
Employees	1,679.71	785.21
	(3,451.24)	(1,269.74)

Notes: Daily stock return is in per cent; its mean and SD (in parentheses) are by firm day, omitting the top and bottom 1% of observed returns. Market capitalisation, turnover and assets are in billion rial. Market capitalisation (from TSE data) is as of March 10, 2014 (TSE, 2014). Total TSE market capitalisation is calculated for firms listed on the TSE as of April 14, 2012, omitting two firms for which data are unavailable. Data from Orbis on turnover, assets and employees are from 2012, and are only available for a subset of firms. The number of firms with available data on turnover and assets is listed in the table; employee data are missing for an additional nine firms (four target firms and five non-target firms).

Factiva measure. The first measure we construct is derived from the Factiva news archive database (Dow Jones, 2019), a business information tool marketed by Dow Jones and Company. Factiva covers around thirty-three thousand global media sources, including most leading newspapers as well as influential multimedia content. We perform a full-text search on Factiva's media archives and retrieve all English-language articles that mention Iran and at least one of the P5+1 countries alongside the word 'sanctions'.[19] The output retrieved from this query is used to create a daily index of sanctions-related news calculated as the total raw article count, standardised within the sample period. Because our goal is to measure the reactions of TSE investors to the events covered in these articles, we assign all articles published over the TSE's weekend (Thursday and Friday) or during holidays to the subsequent trading day.[20]

[19] More specifically, the Boolean query we use is as follows: (Iran) and (US or USA or United States or Russia or China or France or UK or United Kingdom or Germany) and (sanctions).

[20] Baker *et al.* (2016) used a similar approach to construct a monthly measure of economic policy uncertainty by calculating the frequency of articles containing a combination of three words related to uncertainty, drawing on a sample

GDELT measure. Our second event measure is derived from the Global Data on Events, Location and Tone (GDELT) database (GDELT Project, 2018). The GDELT initiative uses machine learning and natural language processing (NLP) tools to create a high-frequency open-source database of political events based on the text of news articles covering more than thirty-seven thousand online news sources in multiple languages (including Persian). A key innovation of GDELT, compared to other news databases, is that it provides information at the event level, where the events are automatically defined via NLP algorithms.[21]

GDELT defines an event as an action undertaken by an actor upon another actor. Actors can be national, subnational (e.g., rebel groups) or transnational (e.g., United Nations). Any event recorded with exactly the same date, subject actor, object actor and event type is treated as a single event and given the same unique event identifier in the GDELT database. This allows a given event to be tracked over multiple articles and days. To measure how much coverage a certain event receives, GDELT reports a count of all source articles mentioning the event and attributes this to the calendar date on which the event first appeared in the news. Each event is also classified into a four-digit category (e.g., '0341 - Express intent to change policy') based on another NLP procedure (Schrodt, 2012).

For our analysis, we extract all political events stored in GDELT for which one of the actors is Iran and the other actor is one of the P5+1 countries. We condition further on events that lie in the two-digit categories that are most likely to contain diplomatic negotiation episodes: '03 - Express intent to cooperate', '04 - Consult' and '05 - Engage in diplomatic negotiations'. In line with our news coverage variable sourced from Factiva, we measure events on an intensive margin using GDELT's count of source articles that mention a certain event, which can be interpreted as a proxy for the event's importance. Again, we assign all events occurring over the TSE's weekend or during holidays to the following trading day.

After calculating the total number of relevant articles for each day, we standardise this measure across all of the days in our sample. Before standardising, we drop all events in February to August 2013, during which time there was a steep drop in the number of articles collected by GDELT, due to technical issues associated with a transition between article collection systems.[22]

Evolution of news coverage. Figure 1 displays the variation in both of our event measures over the course of the main sample period.[23] The top six distinct events identified by each measure, highlighted in the time-series plots in Figure 1, are all related to the sanctions negotiations. Both measures spike at the times when the final framework agreement and final official agreement were completed in Lausanne and Vienna, respectively, when the Geneva deal was reached and when US president Obama made a call to Iran's president, Hassan Rouhani. Our Factiva event measure also clearly identifies a round of negotiations in Vienna that ended in agreement to extend the deadline for a final deal, as well as the release of an International Atomic Energy Agency report in 2012. The GDELT measure instead places more importance on early negotiations in Istanbul

of ten leading US newspapers. In an extensive audit study, they evaluated the performance of this simple text retrieval algorithm by comparing its results with a human-coded index and found the two indices to be highly correlated.

[21] See Manacorda and Tesei (2020) for a recent study that makes extensive use of GDELT data.

[22] For each of the two measures discussed in this section, we also create similar measures that are standardised within the alternative sample period used for the analysis of post-deal events in Subsection 4.3. During this 2016–18 period, when many of the key sanctions-related events did not involve diplomatic negotiations, we also include articles in the GDELT categories '10 - Demand', '11 - Disapprove', '12 - Reject' and '13 - Threaten', along with '16 - Reduce relations' (which includes the subcategory 'Impose embargo, boycott or sanctions').

[23] Variation in news coverage measures for our 2016–18 sample period is discussed in Subsection 4.3.

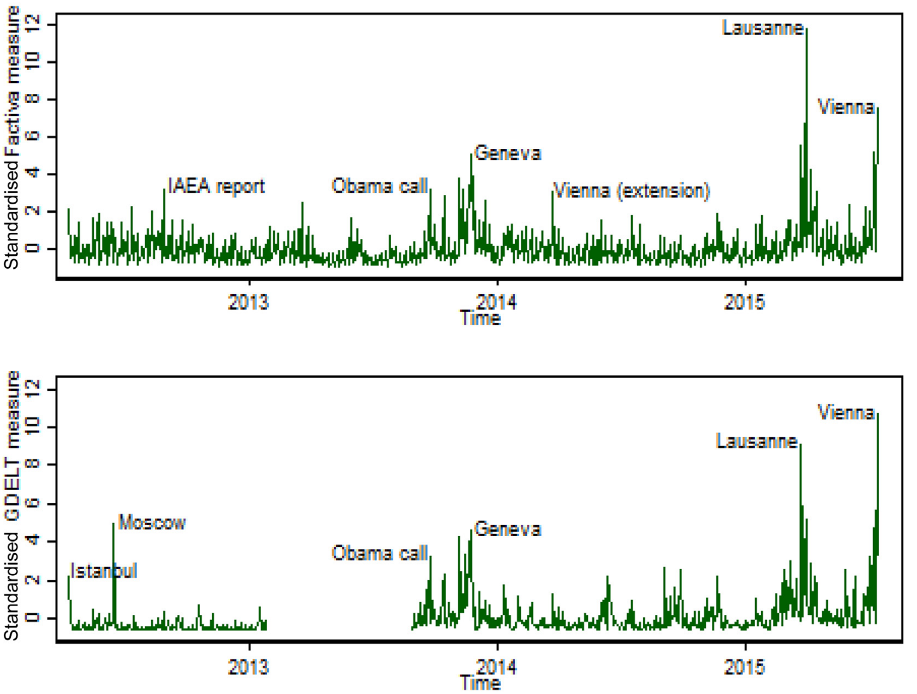

Fig. 1. *Value of Daily News Coverage Measures, April 14, 2012 to July 15, 2015.*
Notes: This figure shows the evolution, over the sample period April 14, 2012 to July 15, 2015 (excluding February to August 2013 for the GDELT measure), of a standardised count of the number of articles on relevant events identified in the Factiva and GDELT datasets, as discussed in Subsection 3.3. The standardised Factiva measure is displayed in the top panel and the standardised GDELT measure is shown in the bottom panel.

and Moscow. This suggests that these measures vary similarly but not identically over time; indeed, the correlation coefficient between the two is 0.72.[24]

3.4. *Other Data*

In order to gain insight into the drivers of sanctions' effects on the non-target portfolio (i.e., their 'bluntness'), we identify the set of Iranian industries subject to sector-specific sanctions. We accomplish this through inspection of the final agreement between Iran and the P5+1 countries (JCPOA, 2015). Annex II of this document ('Sanctions-Related Commitments') provides a list of US and EU sanctions to be removed, along with assurances that these constitute the full set of US and EU sanctions related to Iran's nuclear programme. We check whether each of the measures listed in this annex pertains to a particular Iranian industry or industries, according to the industrial classification used in our paper.

Also, in order to test whether our main results could be driven by an alternative war-related explanation (since diplomatic negotiations may also have changed the likelihood of conflict

[24] The top fifteen days by each measure may be found in Online Appendix Table A1; eight of these are shared across the two measures.

between Iran and other countries), we rely on two additional data sources. First, we extract price data for betting contracts from Intrade, a large online betting market, in order to construct a measure of the probability of military conflict between Iran and Israel or the United States (Intrade, 2014).[25] The betting contract we use for this exercise is specified as 'US and/or Israel execute an overt airstrike against Iran by December 31, 2012'. The contract was to have paid \$10 if an airstrike occurred before December 31, 2012, and zero otherwise. During the trading period, the contract traded in the range of 0 to 100, where 1 point equals \$0.10. We use the contract price to calculate the daily arrival probability of an airstrike that is implied by the contract, under a set of simple assumptions about the conditions in this market.[26] Online Appendix Figure A2 shows the evolution of the contract's daily price and the implied daily arrival probability of an airstrike from the beginning of the sample period (April 14, 2012) through the end of 2012.[27]

Second, we use the Datastream database (Refinitiv, 2019) to extract stock price data for a set of firms for which profit expectations were likely to be sensitive to the odds of war in the Middle East. In particular, we include all publicly traded firms in the Stockholm International Peace Research Institute (SIPRI) 2012 Arms Industry Database (SIPRI, 2016). The SIPRI database identifies the world's hundred largest arms-producing and military service companies in 2012, but since some of these firms are privately owned, we observe stock price data for sixty-six of the companies on this list as of the beginning of our sample period.

4. Empirical Analysis

4.1. *Case Study: Geneva Deal*

To more clearly introduce our empirical strategy and interpretation of our results, we begin by studying the effects of a single major event: the interim agreement negotiated in Geneva in November 2013. We examine this event both because the Geneva deal was perhaps the largest breakthrough in Iran's negotiations with the P5+1 countries, and because the timeline of this agreement (displayed in Figure 2) delivers some useful advantages for our analysis. The final round of negotiations leading to the interim agreement began at 5:00 p.m. on Wednesday, November 20, after the TSE closed at noon on the same day. The TSE was then closed for two days as per the usual structure of the Iranian working week.

Media reports in advance of the meetings indicated uncertainty about whether a deal would be reached, though also some optimism among negotiators that an agreement was possible.[28]

[25] Intrade was an online futures exchange where predictions about events ranging from Oscar nominations to presidential election outcomes were traded as futures contracts. Trading volumes on Intrade reached over one million annually in 2012. Intrade was established in 1999, but in the context of a conflict with the US Commodity Futures Trading Commission, suspended the accounts of its US members in December 2012 and ended trading in early 2013.

[26] Specifically, let $strike_t$ be the number of times an airstrike occurs between day t and December 31, 2012 (which we call day T). On any given day t, investors can buy a contract at price p_t that will be worth ten dollars if $strike_t > 0$ and zero dollars if $strike_t = 0$. Assume that the process is Poisson, investors are risk neutral and do not discount the future, and that the Poisson parameter λ_t prevailing at a given moment is known but can change over time due to information shocks. Then we can back out λ_t from the closing price of the contract, since this price should be equal to $\$10 \times P(strike_t > 0)$, and $P(strike_t > 0) = 1 - e^{-\lambda_t(T-t)}$. Therefore, we can calculate $\lambda_t = (-1/(T-t))\ln(1-(p_t/10))$. We multiply λ_t by a hundred so that it represents a probability in percentage terms.

[27] We do not display data from Intrade after December 22, 2012, since US trading on Intrade was suspended as of December 23. Our regressions using this data similarly use a sample period before the suspension of US trading.

[28] Before the negotiations, a senior US official was reported by Reuters (Wroughton, 2013) to have said that 'I don't know if we will reach an agreement. I think it is quite possible that we can, but there are still tough issues to negotiate.' In the same report, Iran's deputy foreign minister was cited as saying in the Persian-language media that '[t]he expectation is that we will have tough talks, and unless the rights of the Iranian people are guaranteed, an agreement will not be

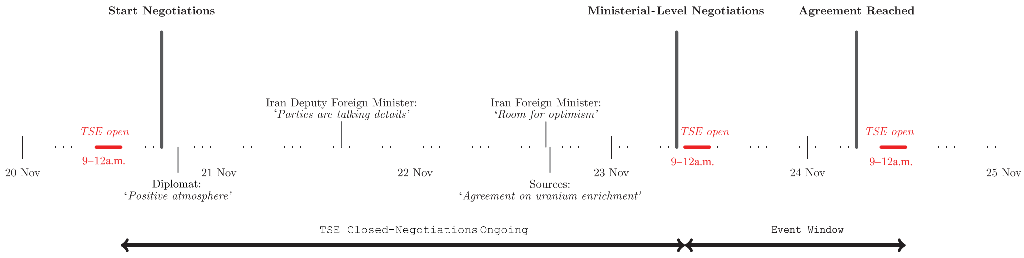

Fig. 2. *Timeline of November 2013 Geneva Event (with Quotes from Nasim Online, 2013 Live Blog).*

During the TSE's weekend break, several encouraging reports emerged about the progress of talks, which continued in Geneva over both Thursday and Friday.[29] When the TSE re-opened at 9:00 a.m. on Saturday, November 23, there was thus a large amount of information on which the market could act. A final agreement was then officially announced at 5:30 a.m. on Sunday, November 24, in advance of the TSE opening again at 9:00 a.m. on the same day.

The positive information shocks before each of these two trading days provide us with an opportunity to use this progress in political negotiations to study the impact of sanctions relief. Specifically, we formulate a difference-in-difference model based on our definitions of the target and non-target portfolios and the timing of news from Geneva. Our baseline model is structured as

$$R_{ijt} = \alpha_i + \phi Geneva_t + \theta Target_i \times Geneva_t + u_{ijt}, \tag{1}$$

where i indexes firms, j represents industries and the time index t is at the day level. Here, R_{ijt} represents raw returns, $Target_i$ is a indicator variable for firms identified as IRGC and/or Setad assets and $Geneva_t$ is a dummy for the two days November 23 and 24, 2013.

We argue that the good news about progress towards a sanctions deal in Geneva should imply greater future benefits for firms that benefit from sanctions relief, because the expected future duration of sanctions falls. If investors hold this belief then the stock prices of these firms should be bid up upon the arrival of this news. Meanwhile, firms unaffected by sanctions removal should not see abnormal returns associated with this event.

The estimated values of θ and ϕ are then informative about the success of sanctions in targeting elite assets within our sample of publicly listed firms. If $\phi + \theta = 0$, the elite assets in our sample escape all of the effects of sanctions, mapping to full bluntness $\beta = 0$ in our conceptual framework. On the other hand, if $\phi = 0$ but the target group interaction coefficient $\theta > 0$, elite assets would be affected without any implied effects for non-target firms. The rejection of both of these hypotheses would provide evidence in favour of the third scenario, in which both the targeted elite group as well as non-elite interests in the sample benefit from sanctions relief. This would imply that the costs of sanctions were shared as a result of the 'blunt instruments' problem.

Baseline results. With these interpretations in mind, we present the results of estimating (1) for a sample period that includes the two days November 23 and 24, 2013 and the sixty trading days

reached' (Wroughton, 2013). According to Agence France Presse (2013), the assessment of an analyst at the International Institute for Strategic Studies was that 'a deal may be possible'.

[29] For example, as shown in Figure 2, an English-language live blog from Nasim Online news agency cited several positive quotes from parties to the negotiations ('positive atmosphere', 'talking details', 'room for optimism'), and reported agreement on the issue of nuclear enrichment, during the TSE's weekend break (Nasim Online, 2013).

Table 2. *Baseline Results—Geneva Deal.*

	(1) Firm FEs	(2) Day-by-day estimates	(3) Firm-day-of-week FEs	(4) Industry controls	(5) Market cap controls	(6) Size data subsample	(7) All size controls
Geneva	0.648		0.731		0.649	0.212	0.221
	(0.277)		(0.286)		(0.279)	(0.356)	(0.353)
Day 1		0.548					
		(0.303)					
Day 2		0.748					
		(0.342)					
Target × Geneva	1.256		1.212	1.306	1.260	1.962	1.927
	(0.364)		(0.383)	(0.415)	(0.373)	(0.466)	(0.483)
Target × day 1		1.415					
		(0.417)					
Target × day 2		1.099					
		(0.445)					
Observations	6,586	6,586	6,580	6,497	6,586	4,017	4,017
Number of firms	128	128	127	127	128	77	77

Notes: This table displays estimated effects of the Geneva deal on returns of target and non-target firms. The dependent variable is daily stock return in per cent. 'Geneva' is defined as the two days Saturday, November 23 and Sunday, November 24, 2013. In column (2), 'day 1' is November 23, 2013 and 'day 2' is November 24, 2013. The sample period is Geneva and the previous sixty trading days. All columns except column (3) include firm fixed effects; column (3) includes firm-day-of-week fixed effects. Column (4) includes interactions of industry dummies with dummies for each day of the sample period. Column (5) includes interactions of de-meaned log market capitalisation as of March 10, 2014 with dummies for each day of the sample period. Column (6) reproduces the baseline estimate in column (1) for the sub-sample of firms (those with available data on turnover, assets and employees in 2012) used in column (7), for purposes of comparison. Column (7) includes interactions of log market capitalisation as of March 10, 2014, and log turnover, log assets and log employees from 2012 (if these data are available from Orbis), all de-meaned, with dummies for each day of the sample period. SEs, clustered by firm, are in parentheses.

beforehand. We use a sixty-day pre-period for our baseline estimates because this is a standard estimation window used in the event study literature.[30] We report SEs clustered by firm, in line with our firm-level definition of 'treatment' as a desired target of sanctions.[31] As shown in column (1) of Table 2, we find that, for non-target firms, the Geneva deal was associated with a daily return that was 0.648 percentage points above the mean return for those firms, with an additional effect of 1.256 percentage points for target firms. Both of these estimates are statistically significant.

We next check whether the effects we estimate are concentrated in either of the two trading days in which the market was hit by positive news about negotiations in Geneva. Column (2) of Table 2 displays the results of regressions in which dummies and interactions for November 23 ('day 1') and November 24 ('day 2') are included separately. The estimated impact is similar across the two days for both the target and non-target groups. Because November 23 and 24 are at the beginning of the trading week, when stock returns might differ systematically from other days, we also rerun our baseline regression substituting firm-day-of-week fixed effects for firm fixed effects, and find very similar results (see column (3)).[32]

Another potential issue with our baseline specification is that firms in the target and non-target groups are somewhat unevenly distributed across industries (see Table 1). Along with the promise of potential changes in the future evolution of sanctions, the Geneva deal also suspended some

[30] Our results are robust to using other windows, including thirty-day or ninety-day periods.

[31] For each of our main tables of results, we also include Online Appendix tables displaying wild bootstrapped p-values (Cameron *et al.*, 2008; Roodman *et al.*, 2019) based on the much stricter strategy of clustering by the fifteen industries in our sample; see footnotes 33, 45 and 63.

[32] Note that because we drop singletons, the number of observations in our regressions sometimes changes when we include different sets of fixed effects.

industry-specific sanctions (on crude oil, petrochemicals, the automotive industry and precious metals). Our results might thus be capturing an asymmetric industry-level effect of the Geneva deal itself rather than the impact of the multilateral sanctions regime as a whole. In Table 2 column (4), we add interactions between industry dummies and dummies for each day in the sample period to our baseline specification, in order to control for industry-specific effects of this event. We find that our estimate of the differential impact of Geneva on firms in the target portfolio is almost unchanged.

We next allow for the possibility that the effects of new information, including good news about sanctions relief, are different for firms with different characteristics (such as firm size). In column (5) of Table 2, we augment the regression of column (1) by controlling for the interaction of the logarithm of market capitalisation (de-meaned) with dummies for each day in our sample period. This has very little effect on our main results.

As noted in our discussion of summary statistics, we also have 2012 data from Orbis on turnover, assets and employees for a subsample of firms. For purposes of comparison, we first reproduce our baseline specification for this subgroup of companies only (see column (6)). While the total effect of Geneva on target firms is similar for this subsample, its impact on non-target firms is smaller than for the full sample (0.212 percentage points) and not statistically significant. In column (7), we add controls interacting log turnover, log assets and log employees (again de-meaned), as well as log market capitalisation, with dummies for each day in the sample period. Our estimates remain very similar to those in column (6), suggesting that these firm characteristics probably do not have an important influence on our full-sample results.[33]

Alternative portfolio definitions. In Table 3, we check the sensitivity of our results to different firm portfolio definitions. We first divide the target portfolio into two parts, consisting of the six firms identified directly from US, EU and UN smart sanctions documents, and the forty-four firms in which targeted entities have a direct ownership stake.[34] Column (1) of Table 3 shows that both sets of firms accrue similar average returns on the Geneva event days. Directly targeted firms have daily returns that are 1.502 percentage points higher than non-target firms, while for other firms in the target portfolio, this difference is 1.227 percentage points. A test of equality between the two coefficients is not rejected. We discuss this finding further in Subsection 4.2.2, with the aim of understanding the relative importance of smart sanctions and more comprehensive sanctions instruments in driving the main results.

We next insert the 172 firms with more tenuous or less certain connections to the IRGC or Setad into the sample as a separate portfolio. As shown in Table 3 column (2), we find that their returns respond to the Geneva deal very similarly to firms in the non-target portfolio. Specifically, the estimated coefficient on the interaction term for this portfolio is small (0.049) and not statistically significant.

As discussed in Subsection 3.2.2, these firms were excluded from our baseline sample for two reasons: they were two or more ownership layers below the seventy-five entities identified in US,

[33] Online Appendix Table A2 displays wild bootstrapped p-values based on clustering by industry for each of the specifications in Table 2. Despite the reduction in power due to the small number of industries, the estimated coefficients on the target interaction term remain statistically significant at conventional levels (along with one p-value of 0.1001). For the estimated effects on the full portfolio of non-target firms in columns (1) to (5), p-values vary between 0.15 and 0.21.

[34] All discussions in the text refer to the full sample of TSE-listed firms as of the beginning of the sample period. In practice, not all firms are present in all regressions, due to firm exit from the TSE and/or thin trading. The number of firms in each regression is noted in the tables, and Table 3 also identifies the number of target and non-target firms.

Table 3. *Alternative Portfolio Definitions—Geneva Deal.*

	(1) Direct targets and assets	(2) Portfolio of other firms	(3) Other sources	(4) Ownership shares	(5) Matched sample	(6) All firms
Geneva	0.648	0.648	0.645	0.757	0.426	0.762
	(0.277)	(0.276)	(0.277)	(0.141)	(0.324)	(0.144)
Target × Geneva		1.256	1.370	2.034	2.043	1.192
		(0.363)	(0.344)	(0.522)	(0.450)	(0.267)
Direct target × Geneva	1.502					
	(0.540)					
Target asset × Geneva	1.227					
	(0.379)					
Other × Geneva		0.049				
		(0.326)				
Observations	6,586	14,842	6,606	13,477	6,920	15,635
Target firms	49	49	50	N/A	23	52
Non-target firms	79	79	79	N/A	108	246
Total number of firms	128	283	129	257	131	298

Notes: This table displays estimated effects of the Geneva deal on returns of various firm portfolios. The dependent variable is daily stock return in per cent. 'Geneva' is defined as the two days Saturday, November 23 and Sunday, November 24, 2013. Column (1) separates the target portfolio into two parts: firms identified directly from smart sanctions documents ('direct target') and firms in which targeted entities have a direct ownership stake ('target asset'). There are six direct target firms and forty-three target asset firms in this regression. Column (2) includes firms dropped from the baseline sample, excluding firms involved in the nuclear programme, as a separate portfolio ('other'); 155 of these other firms are included in this regression. Column (3) defines the target portfolio as the set of firms identified as IRGC or Setad assets in smart sanctions documents, Alfoneh (2010) or Stecklow *et al.* (2013). This adds nineteen additional target firms to the regression, while dropping eighteen firms from the baseline target portfolio. In column (4), the variable 'target' is the observed share of a firm that is owned by firms identified as IRGC or Setad assets in smart sanctions documents, rather than a dummy variable for portfolio membership. In column (5), the sample of firms is selected using coarsened exact matching on industry, market capitalisation, turnover, assets and employees, among firms for which these data are available. In column (6), the sample of firms is widened to include all firms listed on the TSE as of April 14, 2012, and all firms not satisfying the baseline target definition are classified into the non-target portfolio. This also adds three target firms that had been dropped due to involvement in Iran's nuclear programme. The sample period is Geneva and the previous sixty trading days. All columns include firm fixed effects. SEs, clustered by firm, are in parentheses.

EU or UN smart sanctions documents, and/or they were identified as possible assets of the IRGC or Setad by other sources. In columns (3) and (4), we integrate these two types of firms into our sample in turn. We first set aside our information on shareholdings from the TSE, and instead define the target portfolio as the set of firms identified as IRGC or Setad assets by any other source. This still includes the six TSE-listed firms mentioned in US, EU and UN documents, but the remainder of our new target portfolio consists of forty-five additional firms mentioned in Alfoneh (2010) or Stecklow *et al.* (2013), only twenty-five of which were in our baseline target group. Our findings are robust to using this alternative definition, with the estimated differential effect of Geneva on target firms rising slightly to 1.370 percentage points (column (3)).

Next, we instead redefine our baseline specification so that firm-level variation is captured using ownership shares rather than dummy variables for portfolio membership. Specifically, we calculate the total share of ownership of entities identified in US, EU or UN smart sanctions documents in each of their forty-four direct TSE-listed assets, as well as the 144 indirect assets that were excluded from our baseline sample of firms.[35] We also assign ownership shares of one to the six directly targeted firms in our sample, and zero to the eighty-eight non-target firms.

[35] We observe much smaller ownership shares by this measure for most of the indirectly owned firms. While IRGC and Setad shares are below 3% for only nine of the forty-four directly owned assets, this is the case for 133 of the other 144 firms. However, it is important to note that our data on the size of IRGC and Setad ownership shares are likely to be

We then run a regression with a Geneva dummy and its interaction with this ownership variable, along with the usual firm fixed effects. As shown in column (4), we find that firms with no ownership by IRGC and Setad entities have a daily return that is 0.757 percentage points higher on the Geneva event days, and that wholly owned firms see an additional daily return of 2.034 percentage points.

Finally, in columns (5) and (6) of Table 3, we retain our usual definition of targeted entities, but change how we select our sample of firms. In column (5), rather than selecting a non-target portfolio based on our level of certainty about firms' connections to targeted entities, we instead aim to maximise balance between characteristics of firms in the target portfolio and the other firms in the sample. We first take all 325 firms and apply our usual target definition to define a subset of these as 'treated'. We then retain the 193 firms (including 30 target firms and 163 others) with available Orbis data on firm characteristics as of 2012, and employ the method of coarsened exact matching (Iacus *et al.*, 2012) to construct a sample that is balanced by industry, market capitalisation, turnover, assets and employees. Our final sample includes 23 target firms and 111 other firms. Using this matching method, we find an estimated difference due to Geneva of 2.043 percentage points between the average daily returns of target firms and our control group. This estimate is similar to those in columns (6) and (7) of Table 2, which also only included firms with available data on turnover, assets and employees.

In column (6), we instead impose no requirements for sample selection, and simply include all firms listed on the TSE as of the beginning of our sample period. Our target portfolio is defined as usual, but all other firms now serve as the comparison group.[36] This again yields a similar set of results to those in our baseline regression in Table 2 column (1).

Magnitude of effects. These results suggest that both target and non-target firms benefited from sanctions relief, but that the impact on target firms was significantly larger. We now consider whether these effects are economically significant. We first note that the total effect for target group firms of approximately 1.9 percentage points for each day (Table 2 column (1)) is large as compared to the aggregate movements in stock prices usually observed on the TSE. Across our full 2012–15 sample period, the 95th percentile of average daily returns across all firms in the sample is 1.40%, while the 99th percentile is 2.11%.[37] Moreover, the gap in returns between the target and non-target portfolios was also unusually large during the two-day Geneva event, as we show in more detail in a placebo analysis below.

Another useful benchmark is a comparison with the average stock market movements associated with oil price changes (EIA, 2020; OPEC, 2020). We estimate that a one-dollar increase in the price of crude oil is associated with a 0.042 percentage point increase in returns on average among the firms in our sample.[38] This suggests that a forty-five-dollar increase in oil prices would be needed to induce a shift in average returns equivalent to the one-day Geneva effect

very conservative due to our narrow definition of firms owned directly by these two conglomerates (i.e., identification as such by the United States, EU or UN).

[36] Note that our usual target definition now classifies fifty-three firms into the target portfolio, including three firms that had been dropped from the baseline sample due to involvement in Iran's nuclear programme.

[37] Note that the size of our point estimates is realistic when compared to results based on other types of major economy-wide policy shifts, such as changes of government. A useful benchmark here is the study by Snowberg *et al.* (2007). They provided a 2 to 3 percentage point estimate of the impact of a Republican presidency on stock prices across historical elections.

[38] This result is based on the two calendar years immediately prior to our main sample period (2010 and 2011) and is robust across various benchmark prices, as shown in Online Appendix Table A3. The estimated coefficient of 0.042 is from a specification using the West Texas Intermediate price and including firm-quarter-day-of-week fixed effects.

on target firms of 1.9 percentage points. Such daily movements in oil prices are unknown—the largest one-day change during our sample period is approximately eight dollars.

We can also use a back-of-the-envelope calculation to put the 'Geneva effect' on the IRGC and Setad conglomerates into dollar terms. Multiplying total market capitalisation by the sum of the estimates in Table 2 column (1) suggests two-day gains of approximately \$240 million for the six firms directly linked to the IRGC and Setad by the United States, EU and UN, and \$225 million for the shares of the IRGC and Setad in the forty-four other target portfolio firms.

Finally, the impact of Geneva might also be amplified by its impacts on subsequent firm decisions. Although we do not directly observe these effects, we can apply Tobin's q theory of investment to our estimates in order to assess potential consequences for investment behaviour. To do this, we estimate the magnitude of capital adjustment costs using annual panel data from Orbis (Bureau van Dijk, 2020) on realised investment, capital stock and the market-to-book ratio for a subset of the firms in our sample.[39] For the median firm among the thirty-seven target firms for which data are available, our analysis suggests that higher returns to capital due to the two-day Geneva event increased investment growth in 2014 by 0.8 to 2.1 percentage points.[40]

It is important to note that these figures should not be interpreted as capturing the *total* magnitude of the impact of sanctions relief on elite policymakers. In relying on high-frequency stock price movements for identification, we have necessarily excluded the IRGC and Setad conglomerates' substantial (but less readily observable) privately held assets, which could also have been affected by sanctions.[41] But importantly, this exercise does provide evidence that multilateral sanctions against Iran succeeded in inducing some degree of 'income targeting' of the political elite through their publicly listed holdings.

Placebo analysis. The substantial difference between the returns of target and non-target firms is a key driver of the estimated effect of Geneva on elite assets. The statistical significance of this estimate suggests that such a difference between the two portfolios is an unusual occurrence. However, one common concern in stock market event studies is the fact that daily returns do not tend to be normally distributed. Although we have dropped large outliers, we might nonetheless question whether our p-values underestimate the probability that we would observe such a large difference even if the interaction term coefficient $\theta = 0$.

To address this issue, we estimate a series of regressions, each with a dummy for a different two-day period, an interaction of this two-day dummy with our target variable and firm fixed effects. The two-day dummies cover our entire sample period, from April 14–15, 2012 to July 14–15, 2015. We then plot the distribution of the estimated coefficient on the interaction term from each of these 388 regressions in Figure 3, highlighting the 'true' Geneva result. It is apparent

[39] See Online Appendix Table A4. We first run a regression with the ratio of current investment in fixed assets to total fixed assets as of the previous year on the left-hand side, and the previous year's market-to-book ratio on the right-hand side, along with firm and year fixed effects. Another specification adds the ratio of current cashflow to lagged fixed assets as a control variable (Abel and Panageas, 2020). Because of our small sample of firms, we attempt two alternative winsorisation strategies (top and bottom 1% or 5%) and also run regressions including all Iranian firms with available data.

[40] This is in the same range as the forecast effect of US-China tariff policy announcements on 2020 US investment growth in Amiti *et al.* (2020). To reach these results, we use a simplified version of the method suggested by that paper. For each firm, we first use the estimated two-day Geneva effect on market capitalisation to project the change in the market-to-book ratio. We then multiply this by the level of fixed assets in 2013, as well as estimated capital adjustment costs, and divide by 2013 investment in fixed assets. The range of effects cited in the text is based on the smallest and largest estimates of adjustment costs in Online Appendix Table A4.

[41] Stecklow *et al.* (2013) estimated that Setad's real estate holdings alone were worth \$52 billion as of 2008.

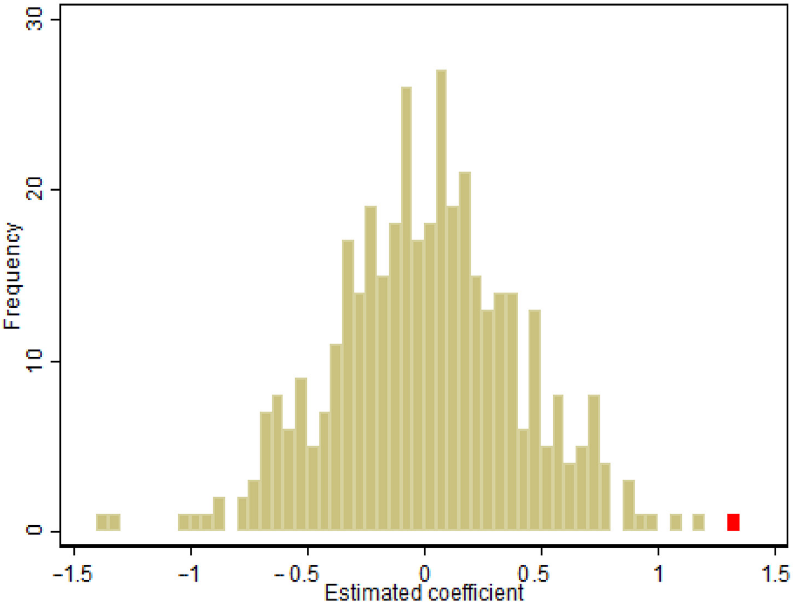

Fig. 3. *Distribution of Estimates across Two-Day Periods (Geneva Deal Highlighted).*
Notes: In this figure, we use data for the period from April 14–15, 2012 until July 14–15, 2015 to estimate a series of models, each with a dummy for a different two-day interval, the interaction of this with our target portfolio dummy and firm fixed effects. Implementation of this procedure for the full set of consecutive two-day intervals corresponds to 388 regressions, and the histogram reports the distribution of estimated coefficients on the interaction term. The two-day Geneva event (November 23–24, 2013) is highlighted (the largest estimated coefficient).

from this figure that our estimate of the differential effect of the Geneva deal on target firms is in the far right tail of this empirical distribution. In other words, over the time period around Geneva, it was highly unusual for the two-day returns of target and non-target firms, conditional on firm fixed effects, to differ from one another to the extent we observe in Table 2.

Daily effects—Geneva sample period. We next take a closer look at day-by-day differences between the returns of target and non-target firms in the period around the Geneva deal. To do this, we return to the sample period from our baseline regression, but add the eight trading days after the end of our two-day Geneva event, so that we may also observe information on the full two weeks after the Geneva shock. We then estimate a series of regressions as described in the previous paragraph, this time instead using one-day dummies for each day in the six TSE trading weeks from Saturday, October 26 to Wednesday, December 4, 2013. We plot the estimated coefficients on the target $\times$ one-day dummy interactions, along with their 95% confidence intervals, in Figure 4.

Throughout the two-week period starting November 23 (to the right of the vertical line in Figure 4), the difference-in-difference coefficient for each day is positive, though not always statistically significant. This implies that the differential effect of Geneva on target firms was not eliminated by movement in the other direction soon after the deal was reached. At the same time, there is a noticeable dip in the coefficients in the weeks preceding the deal, including three days with statistically significant negative estimates. This might have been due to a

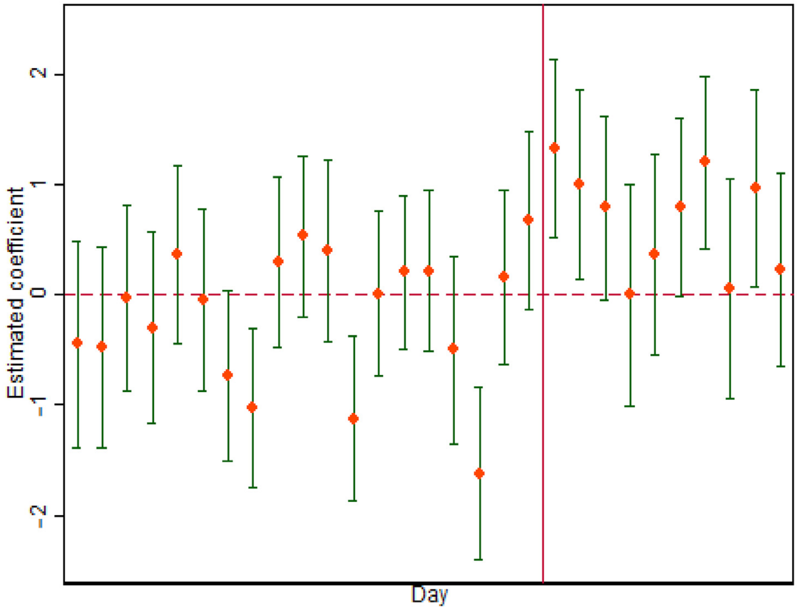

Fig. 4. *Daily Estimated Target-Day Interaction Coefficients around Geneva Deal.*
Notes: In this figure, we use data for the sixty-day estimation window used in Table 2 and the subsequent two weeks to estimate a series of models, each with a dummy for a different day, the interaction of this day dummy with our target portfolio dummy and firm fixed effects. The figure reports the estimated coefficient on the interaction term and its 95% confidence interval for each trading day in the last four weeks of the sixty-day estimation window and the following two weeks. The vertical line is drawn immediately to the left of the first day of the Geneva event (November 23, 2013). SEs are clustered by firm.

temporary surge of pessimism about the potential future incidence of sanctions during the pre-deal negotiations.[42]

However, it is also possible that our positive difference-in-difference estimate of the impact of Geneva is instead driven by mean reversion for target firms after an unrelated negative shock. This uncertainty about interpretation is a limitation of our analysis of a single event. Our next empirical strategy is therefore to use a continuous measure of information shocks, defined over a much longer time period, as outlined in the next section.

Cumulative abnormal returns. As a final exercise, we conduct our event study using a standard approach from the finance literature, by plotting cumulative abnormal returns (CARs) around the event date. We calculate abnormal returns using a standard four-factor model based on Fama and French (1993) and Carhart (1997). Note that we have not used this method as our baseline approach because it purges co-movement between firms' raw returns and market-wide variation, while we use the market-level response to the Geneva event to help measure sanctions' bluntness.

[42] There is also the possibility of early movement in returns due to insider knowledge of the negotiations. We consider this by running similar daily regressions focused on the firms that are most closely connected to elite policymakers. When we separately include an interaction term for the six firms directly linked to the IRGC and Setad (as in Table 3 column (1)), we observe a positive and significant differential return (coefficient 2.131, SE 0.706) for these firms on the last trading day before the event. We also observe a positive and significant coefficient (2.209, SE 0.568) on this day if we define the target variable in terms of ownership shares, as in Table 3 column (4). This might lead us to underestimate the total returns to listed IRGC and Setad assets from progress in Geneva.

However, it is an important robustness check of our finding of differential returns between the target portfolio and other listed firms.

We use the method presented in MacKinlay (1997) to create our CAR figures. We define a forty-day event window, encompassing the twenty trading days preceding the event, as well as twenty days from the first event day onwards. For each firm, we run a regression of raw returns on an intercept and the four factors, using the sixty trading days before the event window as our sample period.[43] We then calculate predicted returns for each firm day during the event window using the estimated coefficients from these regressions. Cumulative abnormal returns for a firm portfolio are defined as the difference between actual and predicted returns, averaged across firms and then cumulated from the beginning of the event window.

Figure 5(a) displays cumulative abnormal returns for our usual target and non-target portfolios, while (b) instead retains all listed firms in the sample, as in the final column of Table 3. The vertical line in the centre of each figure separates the pre-event period (ending on day −1) from the days during and after the event (starting on day 1). In both panels, it is apparent that the CARs accruing to the target portfolio co-move quite closely with those of other firms before the event. During this period, any gaps that appear between the two series are soon closed by movement in the opposite direction. However, following the breakthrough in negotiations in Geneva, the divergence in CARs—favouring target firms—is immediate and persists for the remainder of the event window.

4.2. *News Coverage Analysis*

4.2.1. *Basic news coverage model*

For our main empirical exercise, we extend our sample period to the full range of dates discussed in Subsection 3.1: April 14, 2012 to July 15, 2015.[44] Our period of study begins with the April 2012 Istanbul meeting between Iran and the P5+1 for two main reasons. First, this was the beginning of a diplomatic process with no long-term breakdowns, unlike the several-month gap in talks that preceded Istanbul. Because relevant events occurred throughout the sample, we have sufficient power to identify our parameters of interest using comparisons within shorter sub-periods (quarters). Second, significant changes were made to the multilateral sanctions regime over the several months before April 2012, but the majority of sanctions were in place by the time of the Istanbul meeting. This facilitates an identification strategy based on progress towards full sanctions relief instead of the imposition or removal of particular instruments, allowing us to relate our estimated effects to the sanctions regime as a whole rather than any single sanctions instrument. We end our sample period on July 15, 2015, the day after the final agreement was reached in Vienna.

To capture progress in diplomatic negotiations over the course of this period, we alternately use the daily event measures sourced from the Factiva and GDELT datasets. In both cases, our

[43] We calculate factors—market excess return, small minus big (SMB), high minus low (HML) and daily momentum—according to the variable descriptions provided on Kenneth French's website. Excess return is the value-weighted market return (i.e., returns of all TSE firms weighted by market capitalisation) relative to the one-month Treasury bill rate (sourced from the Federal Reserve Bank of St. Louis FRED dataset (Board of Governors of the Federal Reserve System, 2021)). The value-weighted portfolios used to construct the SMB and HML factors include the 180 of 325 firms for which 2012 book-to-market equity data are available in Orbis. The prior returns used to construct the daily momentum factor are calculated by compounding each firm's daily returns from one year before to one month before each date, excluding the top and bottom 1% of daily returns in the full sample.

[44] As discussed in Subsection 3.3, we exclude February to August 2013 for the GDELT news coverage measure because of a technical issue with GDELT article collection during this period.

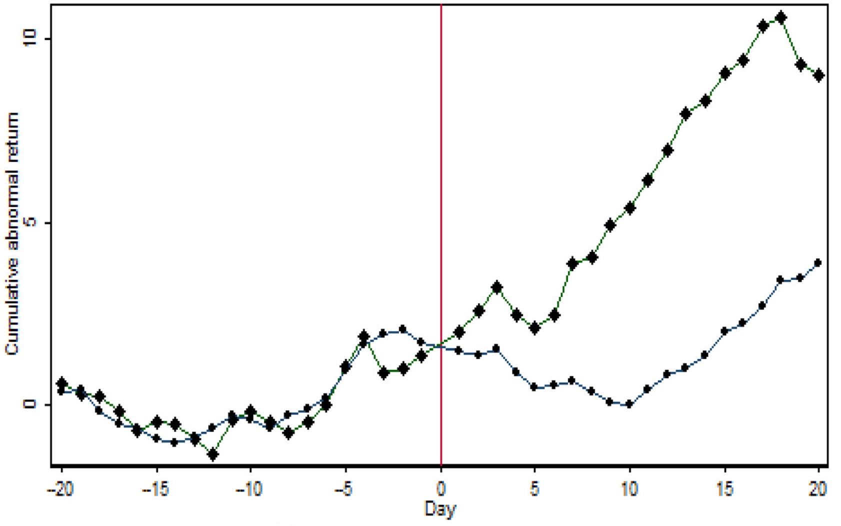

(a) Target and non-target portfolios

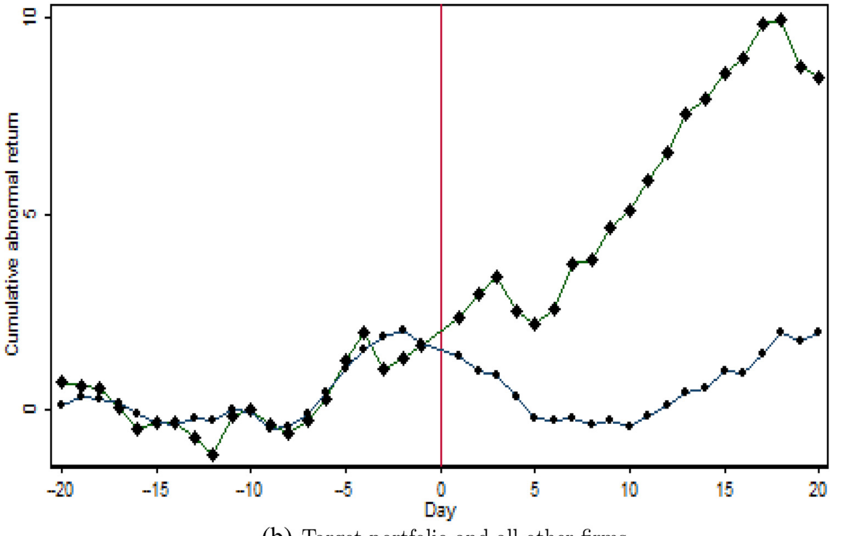

(b) Target portfolio and all other firms

Fig. 5. *Cumulative Abnormal Returns around Geneva Deal.*
Notes: In this figure, we plot cumulative abnormal returns (CARs) over forty days, starting twenty days before the Geneva event. The vertical line is drawn immediately to the left of the first day of the Geneva event (November 23, 2013). Each line represents the CAR for a portfolio of firms, calculated from the difference between actual and predicted returns, averaged across firms. Predicted returns are based on a four-factor model, estimated using the sixty trading days before the event window. In (a), we use the baseline target and non-target portfolio definitions. In (b), we include all firms listed on the TSE as of the beginning of our sample period. Large diamonds represent target firms and small circles represent other firms.

baseline measure of progress is a standardised variable based on the number of news articles about relevant events, as described in Subsection 3.3. We amend the model of (1) by replacing our Geneva dummy with this continuous news coverage variable as

$$R_{ijt} = \alpha_i + \phi Coverage_t + \theta Target_i \times Coverage_t + u_{ijt},$$

where, again, i represents firms, j indexes industries and t indexes days, and the return and target variables are as above.

Column (1) of Table 4 panels A and B displays the results of estimating this specification using the Factiva and GDELT measures, respectively. The Factiva estimates imply that an increase in our news coverage measure by one SD is associated with a return that was 0.055 percentage points above the mean return for non-target firms, and an additional 0.065 percentage points above the mean for target firms. Both of these estimates are statistically significant at the 1% level. Our results using the GDELT measure are similar in magnitude and statistical significance: in this case, when news coverage is higher by one SD, we see average abnormal returns of 0.069 percentage points in the non-target portfolio and 0.064 percentage points in addition to this for the target portfolio.[45]

In column (2), we refine our specification to deal with two potential issues. First, it is possible that our news coverage measure tends to rise over the course of our three-year sample period for reasons other than progress in negotiations (such as better measurement), which could bias our estimates upward if average stock returns also increase. Second, both news coverage measures and stock returns may differ systematically between different days of the week. We thus instead use fixed effects at the firm-quarter-day-of-week level, so that we are, for example, making comparisons between the returns of each firm on Saturdays in the fourth quarter of 2012 with higher and lower values of our coverage variable.[46] This is our preferred specification, and we continue to use firm-quarter-day-of-week fixed effects in all subsequent regressions. Results using the GDELT measure are similar to those in column (1), and for the Factiva measure, our coefficient estimates rise to 0.077 for the coverage variable and 0.092 for its interaction with the target portfolio dummy.

Coverage-based events. The continuous news coverage variables used above are our favoured measures because they are able to capture variation in the relative importance of episodes of diplomatic progress. Indeed, the most important sample days by these measures correspond to the achievement of agreements between Iran and the P5+1 negotiators (see Online Appendix Table A1). However, an alternative possibility is to use the Factiva or GDELT data to define a set of event days, each of which is treated with equal importance in the estimates, i.e., a dummy variable measure. In columns (3) to (7) of Table 4, we present the results of taking such an approach, using a set of increasingly strict definitions of relevant event days.

We begin with a fairly coarse measure in which we assign a value of one to the top 10% of days and zero otherwise. This corresponds (for both Factiva and GDELT) approximately to a dummy for days with a number of relevant news articles at least one SD above the mean. The estimates reported in column (3) show that returns of non-target firms are 0.072 or 0.145 percentage points higher on these event days, and those of target firms are an additional 0.153

[45] The statistical significance of the results in Table 4 is robust to clustering by industry; see Online Appendix Table A5.

[46] 'Quarter' henceforth refers to a unique quarter and year. Our sample ends on July 15, 2015 and thus includes only ten trading days from the third quarter of 2015. In our regressions, we include these ten days in the second quarter of 2015.

Table 4. *Baseline Results—News Coverage.*

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
	Continuous measure						
	Firm FEs	Firm-qtr-day-of-week FEs	Top 10% dummy	2-SD dummy	Spaced 2-SD dummy	Negotiations only dummy	Positive negotiations dummy
Panel A: Factiva measure							
Coverage	0.055	0.077	0.145	0.405	0.439	0.483	0.679
	(0.010)	(0.013)	(0.041)	(0.066)	(0.069)	(0.092)	(0.122)
Target × coverage	0.065	0.092	0.196	0.418	0.363	0.654	0.856
	(0.014)	(0.019)	(0.062)	(0.103)	(0.104)	(0.129)	(0.180)
Observations	75,021	74,775	74,775	74,775	74,094	73,335	72,853
Number of firms	138	136	136	136	136	136	136
Panel B: GDELT measure							
Coverage	0.069	0.065	0.072	0.256	0.235	0.340	0.679
	(0.011)	(0.013)	(0.033)	(0.048)	(0.060)	(0.072)	(0.120)
Target × coverage	0.064	0.066	0.153	0.282	0.303	0.352	0.594
	(0.016)	(0.017)	(0.050)	(0.081)	(0.099)	(0.115)	(0.178)
Observations	62,128	61,843	61,843	61,843	60,763	60,266	59,597
Number of firms	137	136	136	136	136	136	136

Notes: This table displays estimated effects on returns of target and non-target firms from specifications that include a daily measure of news coverage related to diplomatic progress between Iran and the P5+1 countries. The dependent variable is daily stock return in per cent. In all columns, the variable 'coverage' is based on the number of articles on a relevant event identified in the Factiva (panel A) or GDELT (panel B) data. In columns (1) and (2), this is a standardised count. In column (3), this is a dummy for the top 10% of observed values. In column (4), this is a dummy for values that are at least two SDs greater than the mean. Column (5) excludes event days from column (4) that are within one week of larger peaks in news coverage, as explained in Subsection 4.2.1. Column (6) drops event days from column (5) except when these are episodes of direct negotiations between Iran and P5+1 countries. Column (7) excludes event days from column (6) except for episodes where progress in negotiations is apparent from media articles. Column (1) includes firm fixed effects, and columns (2) through (7) include firm-quarter-day-of-week fixed effects. Here, 'quarter' refers to a unique quarter and year. The sample period is from April 14, 2012 to July 15, 2015 (excluding February to August 2013 in panel B). SEs, clustered by firm, are in parentheses.

or 0.196 percentage points larger, using GDELT and Factiva, respectively; all of these estimated coefficients are statistically significant.

We next narrow our definition to encompass only days at least two SDs above the mean. This reduces the number of event days in our sample period to twenty-seven when we use the Factiva data and thirty-one for GDELT. As shown in column (4) of Table 4, our estimates increase substantially in size, to 0.256 (GDELT) or 0.405 (Factiva) percentage points for non-target firms, with a differential effect of 0.282 (GDELT) or 0.418 (Factiva) percentage points for the target portfolio.

A potential issue with this approach is that news coverage may rise in intensity for multiple days around a single event. Progress in diplomatic negotiations may be preceded by speculation about potential meeting outcomes, and followed by news analysis. We therefore sharpen our two-SD measure by pruning our sample in order to 'space' or 'isolate' individual event days.

To do so, we begin from the event day with the largest number of news articles, and drop other event days within one week of this date from the sample. We then proceed to the day with the next highest number of articles (if this has survived the previous step), and repeat the procedure, continuing until we have reached the end of the list of twenty-seven (for Factiva) or thirty-one (for GDELT) days identified above. In terms of the time-series variation visible in Figure 1, this retains the highest peaks within 'clusters' of news coverage as event days, while excluding smaller peaks immediately to the left or right of each of these. As shown in column (5), the results remain similar after the sample is narrowed in this way, even though we have now removed an additional seven event days from our Factiva measure and eleven from GDELT, so that only twenty event days remain in each case.

All of the days identified by this analysis can be linked to sanctions-related events, confirming that the Factiva and GDELT data are very successful in capturing relevant events during our main sample period; a timeline may be found in Online Appendix Table A6.[47] Nonetheless, we again reduce the number of event days by keeping only episodes that can be unambiguously characterised as direct discussions between Iran and P5+1 countries, so as to focus as closely as possible on the returns to diplomatic progress. This leads to the exclusion from the sample of other types of relevant events, such as the release of an IAEA report or a press conference by a key player. We are now left with twelve event days for the Factiva measure and fifteen days for GDELT.[48] Column (6) of Table 4 shows that all of our point estimates increase in size, as compared to column (5), when we impose this restriction.

As discussed in Subsection 2.1, our sample period began with much doubt about whether a deal could ever be reached, and ended with a deal actually agreed. However, some of the individual rounds of negotiations that we have identified as events may not have resulted in substantial forward progress. We therefore estimate a final specification in which we only use events involving a deal being reached, or where diplomatic negotiations were clearly characterised by media reports as leading to progress towards a deal.[49] This further narrows our set of events to seven for Factiva and eight for GDELT; other event days from the previous exercise are dropped

[47] This includes thirty event days in total, because of partial overlap between the two sets of twenty days identified using our two data sources. A regression pooling all thirty events produces similar results to those in Table 4 column (5).

[48] See Online Appendix Table A7 for a timeline of these diplomatic negotiations.

[49] We based this assessment on reports from media organisations based in each of the negotiating countries: Reuters, Associated Press, Agence France Presse, Deutsche Welle, ITAR-TASS, Xinhua, Fars News Agency and Mehr News Agency. See Online Appendix Table A7 for a timeline of headlines from Reuters articles about each set of negotiations.

from the sample. As expected, the estimated coefficients increase when we focus only on these clear cases of progress, as can be seen in column (7).[50]

Online Appendix Figures A3 and A4 reproduce the Geneva day-by-day plot (Figure 4) for the other seven events (excluding Geneva) in Table 4 column (7). The first figure pools all of these events, displaying the daily differences between target and non-target portfolio returns for the days before and after each episode, averaged across the seven event windows. Although there is visible movement in differential returns in advance of these events, there is nonetheless a clear jump on the event day identified by our dummy variable.

The same procedure is applied to each event separately in Online Appendix Figure A4, revealing the idiosyncratic variation underlying these movements in returns. For example, September 2013 saw a multi-day period of openly thawing US-Iran relations, including UN General Assembly speeches by the Iranian and US presidents, culminating in a phone call between the two leaders (panel B). The various news coverage measures in Table 4 each capture different slices of this daily variation: our continuous news coverage measure in column (2) incorporates all of these days, while the dummy variable in column (7) distils the market reaction to the historic phone call.

Magnitude of effects. We now assess the economic significance of the events captured by our news coverage measures. The results for our continuous measure in Table 4 column (2) imply that relative to the median day during our sample period, news coverage at the maximum observed level (i.e., its level on the day of most intensive coverage) is estimated to have had a large impact on returns.[51] The estimated total effect on target firms of 1.44 to 2.04 percentage points (using the GDELT and Factiva measures, respectively) is of a similar magnitude to the result of our Geneva event study above. This is equal to the estimated impact of an oil price rise in the range \$35–\$49.

When we consider the several episodes of positive diplomatic negotiations captured in our specifications in column (7), we find that the average impact on target firms across the seven event days studied in panel A is 1.53 percentage points, and for the eight days in panel B, this is 1.27 percentage points.[52] If we multiply these estimates by the market capitalisation of IRGC and Setad assets and cumulate across days, we find total gains from these seven or eight events of approximately \$640–\$675 million for the firms directly linked to the IRGC or Setad and \$600–\$640 million for their shares in other target portfolio firms. These figures are approximately five times as large as the daily effects we calculated for the Geneva event.[53] This suggests that Geneva was not a one-off phenomenon: IRGC and Setad assets were repeatedly subject to relevant and economically significant shocks during the sample period, as captured by our news coverage measures.

[50] In Online Appendix Tables A8 and A9, we break down the estimates in Table 4 column (5), by running regressions with separate dummy variables for each 'spaced' event identified by Factiva and GDELT, respectively. The tables indicate that the rise in our point estimates between columns (5) and (7) is due to some smaller positive coefficients and some negative coefficients associated with the dropped events. More generally, although our specifications including a wider set of events tend to capture more episodes of progress towards a deal, they are also likely to capture some setbacks.

[51] As shown in Online Appendix Table A1, our two news coverage measures are maximised on two different days, corresponding to the achievement of the framework agreement in April 2015 (Factiva) or the final agreement in July 2015 (GDELT).

[52] As noted during our discussion of the effect of the Geneva deal, the 95th percentile of average daily returns across sample firms during the period of study is 1.40%.

[53] Recall that the figures of \$240 million and \$225 million cited in the previous subsection represented a two-day accumulation.

Of course, because we use daily variation in coverage throughout our sample period, we can also zoom out further to evaluate the effects of a wider range of days. We next use the results of Table 4 column (4), so as to include all days when the intensity of news coverage is at least two SDs above the mean. This augments the numbers in the previous paragraph by around 50% to 100%, suggesting total gains from news on these days of \$1.05–\$1.4 billion for directly linked firms and \$1.0–\$1.3 billion for IRGC and Setad shares in other target firms, depending on whether we use the Factiva or GDELT measure.

4.2.2. *Robustness and heterogeneity*

We next provide a number of robustness checks of our results using the Factiva and GDELT news coverage measures. We begin with an exercise intended to further establish that IRGC and Setad assets were affected by relevant events other than the Geneva deal. Specifically, we rerun our regression using the continuous news coverage measure, from Table 4 column (2), for a period that excludes the entire sample frame used in Subsection 4.1 (the two Geneva event days and the sixty previous trading days). The results, displayed in column (1) of Online Appendix Table A10, are highly reassuring: our estimated coefficients are unsurprisingly somewhat smaller in magnitude, but they remain positive and statistically significant.[54]

Second, as noted above, one key advantage of beginning our sample frame in April 2012 is that the full sanctions regime was essentially in place in its final form by this time, after significant changes over the two preceding years. However, a formal freeze on nuclear-related sanctions by the UN, EU and United States was not in place until November 2013, when it was agreed as part of the Geneva deal. The 'cleanest' test of the effects of diplomatic progress towards comprehensive sanctions relief might therefore exclusively use post-Geneva variation in sanctions-related news.[55] In Online Appendix Table A10 column (2), we find that the relationship between each news measure and returns is again statistically significant for non-target firms and significantly larger for target firms in the period between Geneva and the negotiation of a final agreement.[56]

We also run a further set of regressions based on our earlier robustness checks of the Geneva event study results. We first reproduce the specifications of columns (4) to (7) of Table 2, controlling for daily variation in returns by industry and according to firm characteristics (market capitalisation, turnover, assets and employees). We then consider various alternative portfolio definitions, as in Table 3. These exercises may be found in Online Appendix Tables A10 and A11, using the continuous news coverage measure from Table 4 column (2). We also run the same robustness checks using the positive negotiations dummy from Table 4 column (7): see Online Appendix Tables A12 and A13. We find both sets of results to be robust to all of these alternative specifications.

Potential role of smart sanctions. We now briefly explore the role of 'smart sanctions' instruments, as opposed to more comprehensive instruments such as restrictions on Iran's entire

[54] The results in column (7) of Table 4 are also not driven by any particular event, as they are robust to dropping any one event day at a time.

[55] Note that the United States continued to impose economic penalties on non-Iranian companies violating the existing sanctions regime, for example in December 2013 (see Online Appendix Table A6).

[56] Recall that to construct the target and non-target portfolios, we use shareholder data from April 14, 2012, the first day of our sample period. To explore the implications of potential changes in ownership over time, we re-estimate the regression in Online Appendix Table A10 column (2) using target and non-target portfolios that are instead constructed via TSE shareholder data (TSE, 2017) from November 20, 2013 (the final trading day before the Geneva deal). Results using the Factiva measure remain of similar magnitude and statistical significance, but in the case of the GDELT measure, the estimated coefficient on the interaction term is smaller by half and has a p-value of 0.19.

financial sector, in driving the targeting effects we observe.[57] Although most sanctions instruments were already in place before the beginning of our main sample period, the imposition of smart sanctions on Setad assets by the US Treasury in June 2013 (see Subsection 2.2.2) may have had an important effect. In Online Appendix Table A14, we look at sensitivity to sanctions news measures separately by conglomerate, before and after mid-2013. While returns of IRGC assets are more responsive to sanctions-related news than returns of non-target firms throughout the sample period, analogous effects for Setad-owned firms are apparent only after smart sanctions were imposed. This is suggestive evidence of the efficacy of these particular sanctions instruments.[58]

We next compare the target firms that were directly subject to smart sanctions with the other IRGC and Setad assets in the target portfolio. As already reported in column (1) of Table 3 and Online Appendix Tables A11 and A13, the estimated impact of progress towards a diplomatic agreement is generally very similar for each of these subgroups.[59] The only exception is the finding in Online Appendix Table A13 panel B, where the effect of the episodes of diplomatic progress identified by GDELT is estimated to be significantly larger for directly targeted firms. But even in this specification, IRGC and Setad assets that are not subject to smart sanctions still see significantly larger returns than non-target firms.

Rather than providing evidence against smart sanctions' effectiveness, this finding might instead suggest that their impacts spill over to other related firms. The study of sanctions on Russia by Ahn and Ludema (2020) notes the phenomenon of 'de-risking', in which potential business partners shy away from transactions with even an indirect connection to a firm subject to smart sanctions. If we measure the strength of this connection in terms of the ownership share of parents subject to IRGC or Setad-related smart sanctions, our earlier results in column (4) of Table 3 and Online Appendix Tables A11 and A13 suggest that firms with stronger connections to these parents are significantly more sensitive to sanctions news.[60] Overall, these exercises suggest that smart sanctions might have been an important driver of our main results, but the evidence is not conclusive.

Bluntness of sanctions. A related question is the reason behind the responsiveness of other firms to sanctions-related news. Our results suggest that despite having a larger impact on assets of the IRGC and Setad, the sanctions are nonetheless 'blunt', as they also affect TSE-listed firms outside the target portfolio. That is, in terms of our conceptual framework, we have found evidence that $\beta < 1$.

We now consider a possible mechanism behind the bluntness of sanctions in the context of the firms we study. As discussed in Subsection 1.2, this might be a direct consequence of elements of the sanctions themselves, or occur indirectly due to spillovers from the parts of the Iranian economy that were subject to sanctions. Here, we test for the direct incidence of sanctions on firms outside the target portfolio, to check if at least some of the 'blunt instruments' problem is a direct effect of the instruments used in the Iran sanctions regime.

[57] Note that even in the absence of smart sanctions, more comprehensive instruments could generate the results we observe if their impacts on IRGC and Setad assets were stronger than their effects on other firms.

[58] An event study using the date of announcement of the Setad sanctions does not turn up significant related short-run movements in returns. This might be because the sanctions were anticipated by investors, but unfortunately, we are not aware of whether information about these sanctions became available before they were officially announced.

[59] For the continuous news coverage measures, this conclusion is based on the full sample period, but also holds when we consider only the period after the United States imposed smart sanctions on Setad.

[60] These specifications included directly targeted and non-target firms, but the results also hold when we limit the sample to firms with indirect ownership by parents subject to IRGC or Setad-related smart sanctions.

To do so, we take advantage of the fact that many of the comprehensive instruments used in the sanctions regime were sector specific, such as US sanctions on Iran's automotive industry. In particular, we check whether non-target firms in Iranian industries subject to sanctions are more responsive to progress towards a diplomatic agreement. As discussed in Subsection 3.4, we identify these industries through inspection of the final deal between Iran and the P5+1 countries (JCPOA, 2015), which provides a full list of the sanctions to be removed.[61]

In Online Appendix Table A15, we reproduce each of our main specifications—the Geneva event study regression from Table 2 column (1), and our news coverage regressions from Table 4 columns (2) and (7) – but replace the target portfolio dummy in the interaction term with a dummy for industries subject to sanctions. Our sample consists only of the firms in the non-target portfolio. Because the variable identifying affected firms is now determined at the industry level, we report wild bootstrapped *p*-values based on clustering by industry.

Our estimates of the coefficient on the interaction term are consistently positive and large, suggesting that firms in sanctioned industries are substantially more responsive to sanctions-related news. Moreover, despite limited statistical power due to the coarseness of our industrial classification, these results are statistically significant in two of our five regressions (columns (4) and (5)), and the relevant *p*-values fall between 0.1 and 0.2 in the other three specifications. So although our previous results indicate that target firms experience larger effects than non-target firms—even within the same industry—we also observe evidence of direct impacts of sanctions on non-target firms in sanctioned industries.[62]

4.3. *2016 US Election and Subsequent Events*

We next consider a relevant and unexpected event outside our main sample period: the 2016 US election. When the TSE opened on November 9, 2016, it was apparent that Donald Trump had very likely been elected US president, a surprising shift from the probabilities prevailing at the time of the market's closure on the previous day.[63] Because Trump had signalled his opposition to the Iran deal during the 2016 election campaign, while Hillary Clinton was a supporter of the agreement, this event provides us with an opportunity to study a sudden, unanticipated rise in the probability of sanctions being reimposed by at least one of the deal's signatories.

As shown in column (1) of Table 5, the TSE's response to this shock is a mirror image of the estimated impact of the Geneva deal in column (1) of Table 2 (from which we reproduce the same specification, including a sixty-day pre-period and firm fixed effects). Average returns for non-target firms are 1.667 percentage points below their mean, while the abnormal return

[61] According to the industrial classification used in our paper, the sanctioned Iranian industries are mining (which includes oil extraction), food products, wood/paper/textiles, refined petroleum, chemicals (which includes petrochemicals), motor vehicles, transportation/telecom and finance. Of the eighty-eight firms in our non-target portfolio, thirty-eight belong to these industries.

[62] Notably, the gap between the estimated effects for firms in targeted and non-targeted industries is largest when we focus on the Geneva deal (column (1)). Given our lack of statistical power throughout Online Appendix Table A15, this finding should be treated cautiously. But it is consistent with the observation that the Geneva deal may have been especially important for firms in sanctioned industries, because it included a clause suspending some industry-specific sanctions (as discussed in Subsection 4.1).

[63] See Silver (2016) for a discussion of expectations and predictions ahead of the 2016 US election. Silver's FiveThirtyEight model had the highest (29%) probability of Trump winning amongst a range of polling-based predictions, typically in the 1% to 15% band. In addition, Silver notes that betting markets put Trump's odds at 18% on the eve of the election.

Table 5. *Post-Deal Events—Potential Restoration of US Sanctions.*

	(1) US election Nov 2016	(2) Factiva measure post-election	(3) GDELT measure post-election	(4)
Election	−1.667			
	(0.305)			
Target × election	−1.296			
	(0.434)			
Coverage		0.040	0.051	
		(0.017)	(0.016)	
Target × coverage		−0.031	−0.036	
		(0.021)	(0.023)	
US coverage				0.015
				(0.016)
Target × US coverage				−0.007
				(0.023)
P4+1 coverage				0.044
				(0.019)
Target × P4+1 coverage				−0.035
				(0.027)
Observations	6,346	36,484	36,484	36,484
Number of firms	120	123	123	123

Notes: This table displays estimated effects of the 2016 US election and subsequent sanctions-related events on returns of target and non-target firms. 'Election' is defined as the event day Wednesday, November 9, 2016. 'Coverage' is a standardised count of the number of articles on a relevant event identified in the Factiva (column (2)) or GDELT (columns (3) and (4)) data. 'US coverage' is a standardised count of the number of articles on a relevant event in the GDELT data in which the United States is the country other than Iran identified by the GDELT dataset, while 'P4+1 coverage' is an analogous measure for the other members of the P5+1 group. The dependent variable is daily stock return in per cent. In column (1), the sample period is the election event day and the previous sixty trading days. In columns (2), (3) and (4), the sample period is from November 12, 2016 to May 9, 2018. Column (1) includes firm fixed effects and columns (2), (3) and (4) include firm-quarter-day-of-week fixed effects, where 'quarter' refers to a unique quarter and year. SEs, clustered by firm, are in parentheses.

for target firms is even more negative: an additional 1.296 percentage points below average.[64] Figure 6 presents daily event study evidence similar to that in Figure 4, with the day of the election standing out as displaying significant negative effects for target group firms relative to the non-target portfolio. This response to the US election surprise is consistent with our earlier conclusions regarding the targeting of nuclear-related sanctions across listed firms in Iran.[65]

Changes in probability of sanctions relief. Because of the availability of time-varying estimates of the probability of Trump's election, we can use this event to generate estimates of changes in the probability of sanctions relief due to other events we have studied, and the Geneva deal in particular. While this exercise requires some strong assumptions, it allows us to assess whether the implied movements in probability are broadly reasonable, and to make further conjectures about the economic magnitude of our estimates.

According to betting markets, the probability of a Trump victory rose by approximately 80 percentage points between the TSE's closing times on November 8 and 9 (see footnote 62).

[64] Online Appendix Table A16 shows that the statistical significance of the results in Table 5 is robust to clustering by industry. Online Appendix Tables A17 and A18 present robustness checks of the result in Table 5 column (1), based on the specifications in Table 2 columns (3) to (7) and Table 3.

[65] Of course, the election of Trump rather than Clinton represents a choice between a whole suite of likely policies, and we cannot separate the effect of a possible shift in US sanctions policy from impacts of other potential changes—though we consider one key factor, conflict risk, in Section 5.

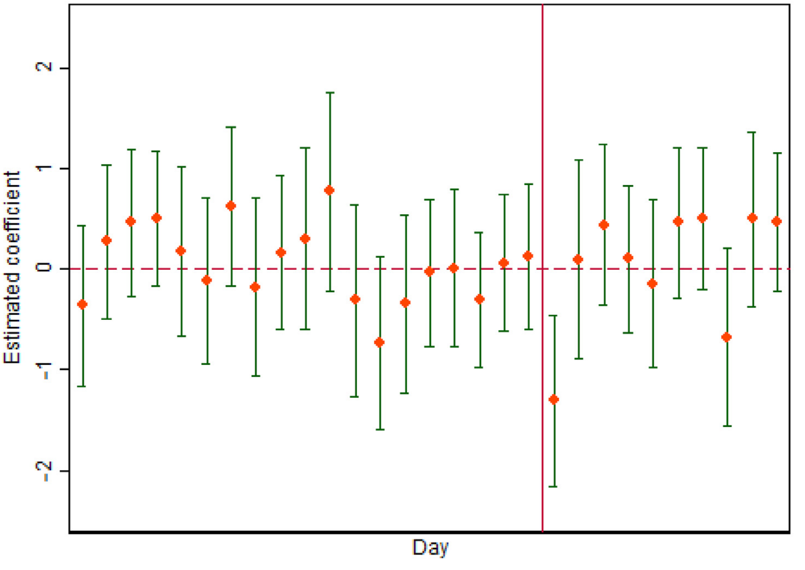

Fig. 6. *Daily Estimated Target-Day Interaction Coefficients around 2016 US Election.*
Notes: In this figure, we use data for a period including the sixty-day estimation window used in Table 5 column (1) and the subsequent two weeks to estimate a series of models, each with a dummy for a different day, the interaction of this day dummy with our target portfolio dummy and firm fixed effects. The figure reports the estimated coefficient on the interaction term and its 95% confidence interval for each trading day in the last four weeks of the sixty-day estimation window and the following two weeks. The vertical line is drawn immediately to the left of the US election event (November 9, 2016). SEs are clustered by firm.

We make the simple assumptions that Trump's defeat of Clinton increased the probability of unilateral restoration of sanctions by 50 percentage points, that US sanctions constitute 50% of the economic costs of multilateral sanctions, and that sanctions imposition and relief can be treated as symmetric. Then we can reach an estimate of the impact of full multilateral sanctions relief on returns by scaling the election estimate by a factor of five, to approximately 15 percentage points for the target portfolio.[66]

This is four times the size of the two-day Geneva effect on target firms (3.8 percentage points), calculated from Table 2 column (1). The implication is that the successful conclusion of the Geneva deal raised the probability of multilateral sanctions relief by around 25 percentage points, a realistic possibility. If our estimates of rises in the value of IRGC and Setad assets are scaled accordingly, we arrive at an estimated return from full sanctions relief of approximately $1.8 billion.

Post-election events. We also attempt to investigate the responses of target and non-target firms to later sanctions-related events, using our Factiva and GDELT news coverage measures. We study a sample period starting on the first trading day after the election event (November 12, 2016) and ending with the US withdrawal from the Iran deal (May 9, 2018). In columns (2) and (3) of Table 5, we see that the estimated coefficient for non-target firms is positive; as shown in

[66] This is the sum of the two estimates in Table 5 column (1), divided by the change in probability of restoration of US sanctions (0.8×0.5), divided by the share of US sanctions in total economic costs (0.5).

column (4), this is driven by news involving the non-US signatories to the sanctions agreement, who maintained support for the deal despite US opposition.[67] However, there is no longer a differential effect on target firms: the response of their stocks to these events is not statistically different from that of the non-target portfolio.

In contrast to the previous result, this suggests that senders' sanctions policy changes drifted 'off target' at some point after the election. However, the post-2015 variation in our news coverage measures does not isolate shocks related to the multilateral nuclear deal as cleanly as before. As seen in Online Appendix Table A20, which lists the top fifteen event days for each measure during this period, we now also capture events such as Iranian missile tests and a confrontation at sea between the United States and Iran. We therefore cannot interpret the findings in Table 5 columns (2) to (4) as relating to the expected incidence of sanctions as confidently as in Subsection 4.2.[68]

5. Alternative Interpretation: Conflict Risk

We now look more deeply into the assumption that the content of our event dummies and news coverage measures consists of 'clean' information that is predominantly about sanctions. Progress in diplomatic negotiations may have had other implications for the foreign policy environment faced by Iran, specifically the probability of military conflict between Iran and other countries. We perform two exercises to assess the potential implications of this issue for the interpretation of our results, focusing alternatively on Iranian and non-Iranian firms.

5.1. *TSE Firm Returns and Conflict Risk*

Our first exercise looks at the relative sensitivity of target and non-target firms to the risk of direct conflict. In order to measure the probability of conflict, we use political betting market prices for a contract regarding the likelihood that the United States and/or Israel would 'execute an overt airstrike against Iran by December 31, 2012', as discussed in Subsection 3.4. We then estimate the relationship of the stock returns of firms in each portfolio with the change in the daily arrival probability of an airstrike implied by these contract prices. To do this, we apply our specification from Table 4 column (2), with the day-to-day change in the probability of conflict replacing our news coverage measure, using data from the first day of our sample period to the end of 2012.[69]

The results indicate that increases in conflict probability, inferred from rises in contract prices, negatively affect stock returns for firms in our sample. In Table 6 column (1), we find that an increase in the daily arrival probability of an airstrike that is higher by 0.1 percentage points (which is approximately three times the SD of this variable) is associated with a statistically significant fall in returns of 0.158 percentage points. However, there is no differential sensitivity

[67] To define separate regressors for news involving the United States and other countries in the P5+1 group, we use the fact that in the GDELT data each event is identified with only one country other than Iran.

[68] In Online Appendix Table A19, we instead study a set of events based on the requirement for the US president to renew a waiver on most US sanctions every 120 days. This waiver, which was first signed by Barack Obama for the implementation of the 2015 Iran deal, was renewed by Donald Trump in May and September 2017 and January 2018, but not in May 2018, when Trump instead withdrew from the deal. Unfortunately, news reports suggest that each of these decisions was expected in advance. Nonetheless, we show the results of event studies based on Trump's three renewals (columns (1) to (4)) and the US withdrawal (column (5)). We observe impacts of both sets of events on non-target firms—positive when the waiver was renewed and negative when it was not—but again see no differential impact on target firms.

[69] Ideally, we would also like to use information from similar contracts later in our sample period. However, as noted earlier, Intrade suspended its US accounts in late December 2012, and all of its activities in early 2013.

Table 6. *Alternative Explanation—Conflict Risk.*

	(1)	(2)	(3)	(4)	(5)
	Intrade probability	Arms industry firms			
		Factiva measure		GDELT measure	
Change in arrival probability	−1.580				
	(0.606)				
Target × change in arrival probability	−0.783				
	(0.845)				
Coverage		−0.00005	−0.012	0.008	0.009
		(0.007)	(0.014)	(0.006)	(0.014)
Sensitivity rank × coverage			0.0004		−0.00003
			(0.0004)		(0.0004)
Observations	13,885	50,586	50,586	41,514	41,514
Number of firms	133	66	66	66	66

Notes: This table displays the results of regressions exploring the possibility that the main results are driven by changes in conflict risk. Column (1) displays correlations between returns of target and non-target firms and the day-to-day change in the arrival probability of a US or Israeli airstrike against Iran, calculated from the price of the Intrade contract ‘US and/or Israel to execute an overt airstrike against Iran by December 31, 2012’. Columns (2) to (5) display estimated effects on returns of non-Iranian firms in the arms industry from specifications that include a daily measure of news coverage related to diplomatic progress between Iran and the P5+1 countries. The dependent variable is daily stock return in per cent. ‘Change in arrival probability’ represents the change in the arrival probability of a US or Israeli airstrike against Iran, during the same day, in per cent. ‘Coverage’ is a standardised count of the number of articles on a relevant event identified in the Factiva (columns (2) and (3)) or GDELT (columns (4) and (5)) data. ‘Sensitivity rank’ is the rank of each firm based on the responsiveness of its stock return to the change in the airstrike arrival probability in the first quarter of 2012, as discussed in Subsection 5.2. All columns include firm-quarter-day-of-week fixed effects. Here, ‘quarter’ refers to a unique quarter and year. In column (1), the sample period is from April 14, 2012 to December 22, 2012, and in columns (2) to (5), it is from April 16, 2012 to July 15, 2015. SEs, clustered by firm, are in parentheses.

between the two portfolios, since the estimated coefficient on the interaction of our target group dummy with the change in arrival probability is statistically insignificant. In short, while conflict risk appears to have a negative effect on firm returns across the TSE, there is no indication that the target group is more economically vulnerable to this risk.

5.2. *Information Content of Sanctions-Related News Measures*

In our second exercise, we examine the possibility that our news coverage measures contain confounding information about the probability of direct military conflict. The ideal sample for this exercise would be a group of firms that are ‘war sensitive’ but not exposed to the effects of sanctions relief. Our approach to approximating this ‘ideal set’ of firms is to identify a group of companies on stock markets outside Iran who are sensitive to conflict risk (benefiting from increased odds of conflict) but, as non-Iranian firms, are not directly exposed to the local effects of sanctions.

As discussed in Subsection 3.4, we use a group of sixty-six firms from the 2012 SIPRI Arms Industry Database, which identifies the world’s largest arms-producing and military service companies. We first confirm that this portfolio is sensitive to the odds of Iran-related conflict, since these firms’ daily stock returns are strongly related to the change in the daily arrival probability of an airstrike implied by Intrade contract prices.[70] This allows us to use them as a placebo group to investigate whether conflict risk and sanctions news are confounded.

[70] More specifically, we find that a specification with daily return as the dependent variable and the change in the daily arrival probability of an airstrike implied by Intrade contract prices as the regressor, along with firm-quarter-day-of-week fixed effects, yields an estimated coefficient of 2.813 and a SE of 0.372.

In columns (2) and (4) of Table 6, we regress the daily returns of arms industry firms on each of our news coverage measures and firm-quarter-day-of-week fixed effects, as in the specification in column (2) of Table 4, using data from the full 2012–15 sample period. The estimated coefficient on the coverage variable is statistically insignificant in each case. We then refine our exercise to check whether firms that are more sensitive to the odds of conflict between Iran and other countries respond differently to these news shocks. For each of the firms in the sample, we run a regression of daily returns in the quarter before the sample period (January to March 2012) on changes in airstrike arrival probability. We then rank the firms according to the size of the estimated coefficient on airstrike arrival probability (from most positive to most negative) and add the interaction between this variable and our sanctions news measure to our regressions. This again yields statistically insignificant results.

In Online Appendix Table A21 panel A, we repeat this exercise using a Geneva event dummy rather than our news coverage measures, and similarly find no evidence that the returns of our sample of arms industry firms were affected by this sanctions-related event. In the same table, we show that the 2016 US election was followed by a surge in the stock market returns of these firms. Gains in defence stocks after Trump's victory were well documented in media reports at the time, which cited expected rises in military spending under Trump (e.g., Heath, 2016; Marcial, 2016). However, we do not find evidence that this post-election bump was linked to heightened risk of conflict with Iran in particular, as firms ranking more highly in the sensitivity measure introduced above did not experience differential gains.[71]

6. Conclusion

In this study, we test a central plank of international sanctions policy: the efficacy of the targeting of elite decision makers within a sanctioned country. We consider the case of Iran, examining the response of its stock market to information, indicating progress in diplomatic negotiations towards an agreement on sanctions removal. We find evidence that listed firms owned by key groups within Iran's political system were differentially sensitive to news about potential sanctions relief leading up to the final deal. Importantly, this sensitivity suggests that political elites within Iran faced a tangible economic incentive to negotiate the removal of sanctions.

The change in US policy towards Iran after the 2016 election raises the question of whether the response of Iranian firms to sanctions-related news has continued to follow a pattern consistent with successful targeting. Our analysis indicates that, while target firms responded more sharply to Donald Trump's election, such differential returns are not evident for later events leading up to the 2018 US withdrawal from the Iran deal. However, news about sanctions policy is more difficult to interpret from 2016 onwards relative to the 2012–15 period, when members of the P5+1 presented a unified front and news coverage was a reasonable proxy for forward progress in negotiations.

It should also be noted that, while our focus on listed firms has made our study possible by allowing us to use high-frequency news shocks for identification, there are other important dimensions of the Iranian economy that we do not measure here. First, our findings exclude the costs from sanctions that were incurred by households. These are likely to have been considerable:

[71] In contrast, the interaction of each of our 2016–18 news coverage measures with this firm-specific measure of sensitivity to Iran-related conflict does have a statistically significant estimated coefficient (see Online Appendix Table A21 panel B). This again suggests that the 2016–18 versions of our Factiva and GDELT measures may be capturing additional information other than sanctions news.

in 2012, when the intensity of the sanctions regime reached its peak, Iranian GDP declined by 7.4%, while inflation exceeded 25% (World Bank, 2019). Our results for listed firms therefore need to be put into the context of other likely economic and social costs. Second, we do not observe assets of the IRGC and Setad that are not listed on the TSE. Our results do, however, provide evidence that sanctions were 'on target' in the sense that sources of income for elite policymakers were positively affected by their removal. The evidence we present here therefore indicates that a 'complete policy failure' scenario—where Iran's political elites fully escaped direct negative income effects from sanctions—can be ruled out.

University of Warwick, UK
University of Ottawa, Canada
Deloitte UK, UK
Queen Mary University of London, UK & LICOS Centre for Institutions & Economic Performance, Belgium

Additional Supporting Information may be found in the online version of this article:

Online Appendix
Replication Package

References

Abel, A.B. and Panageas, S. (2020). 'An analytic framework for interpreting investment regressions in the presence of financial constraints', Working Paper 26898, National Bureau of Economic Research.
Acemoglu, D., Hassan, T.A. and Tahoun, A. (2018). 'The power of the street: Evidence from Egypt's Arab Spring', *Review of Financial Studies*, vol. 31(1), pp. 1–42.
Agence France Presse. (2013). 'Iran says next nuclear talks will be "difficult"', 17 November, retrieved from Factiva database, document AFPR000020131117e9bh0083v (last accessed: 18 November 2019).
Ahn, D.P. and Ludema, R.D. (2019). 'Measuring smartness: The economic impact of targeted sanctions against Russia', in (T. Besedes and V. Nitsch, eds.), *Disrupted Economic Relationships: Disasters, Sanctions, Dissolutions*, pp. 131–54, Cambridge, MA: MIT Press.
Ahn, D.P. and Ludema, R.D. (2020). 'The sword and the shield: The economics of targeted sanctions', *European Economic Review*, vol. 130, p. 103587.
Alfoneh, A. (2010). 'The Revolutionary Guards' looting of Iran's economy', *Middle Eastern Outlook*, no. 3, Washington, DC: American Enterprise Institute.
Allen, S.H. (2008). 'The domestic political costs of economic sanctions', *Journal of Conflict Resolution*, vol. 52(6), pp. 916–44.
Amiti, M., Kong, S.H. and Weinstein, D.E. (2020). 'The effect of the U.S.-China trade war on U.S. investment', Working Paper 27114, National Bureau of Economic Research.
Baker, S.R., Bloom, N. and Davis, S.J. (2016). 'Measuring economic policy uncertainty', *Quarterly Journal of Economics*, vol. 131(4), pp. 1593–636.
BBC. (2012). 'Iran nuclear talks in Istanbul "constructive"', http://www.bbc.com/news/world-middle-east-17716241.
Berman, E., Shapiro, J.N. and Felter, J.H. (2011). 'Can hearts and minds be bought? The economics of counterinsurgency in Iraq', *Journal of Political Economy*, vol. 119(4), pp. 766–819.
Board of Governors of the Federal Reserve System. (2021). 'Market yield on U.S. Treasury securities at 1-month constant maturity [dataset]', https://fred.stlouisfed.org/series/DGS1MO (last accessed: 15 June 2021).
Borger, J. (2015). 'Iran nuclear negotiations enter final day as deadlock persists', *Guardian*, 31 March.
Bureau van Dijk. (2020). 'Orbis [dataset]', https://www.bvdinfo.com/en-gb/our-products/data/international/orbis (last accessed: 27 June 2020).
Cameron, A.C., Gelbach, J.B. and Miller, D.L. (2008). 'Bootstrap-based improvements for inference with clustered errors', *Review of Economics and Statistics*, vol. 90(3), pp. 414–27.
Carhart, M.M. (1997). 'On persistence in mutual fund performance', *Journal of Finance*, vol. 52(1), pp. 57–82.
Čihák, M., Demirgüç-Kunt, A., Feyen, E. and Levine, R. (2012). 'Benchmarking financial systems around the world', Policy Research Paper 6175, World Bank.
Coulomb, R. and Sangnier, M. (2014). 'The impact of political majorities on firm value: Do electoral promises or friendship connections matter?', *Journal of Public Economics*, vol. 115, pp. 158–70.

Dell, M. and Querubin, P. (2018). 'Nation building through foreign intervention: Evidence from discontinuities in military strategies', *Quarterly Journal of Economics*, vol. 133(2), pp. 701–64.
DellaVigna, S. and La Ferrara, E. (2010). 'Detecting illegal arms trade', *American Economic Journal: Economic Policy*, vol. 2(4), pp. 26–57.
Dow Jones. (2019). 'Factiva [dataset]', https://www.dowjones.com/professional/factiva/ (last accessed: 28 January 2019).
Dube, A., Kaplan, E. and Naidu, S. (2011). 'Coups, corporations and classified information', *Quarterly Journal of Economics*, vol. 126(3), pp. 1375–409.
Dube, O. and Naidu, S. (2015). 'Bases, bullets and ballots: The effect of US military aid on political conflict in Colombia', *Journal of Politics*, vol. 77(1), pp. 249–67.
Durante, R. and Zhuravskaya, E. (2018). 'Attack when the world is not watching: U.S. media and the Israeli-Palestinian conflict', *Journal of Political Economy*, vol. 126(3), pp. 1085–133.
Eaton, J. and Engers, M. (1992). 'Sanctions', *Journal of Political Economy*, vol. 100(5), pp. 899–928.
Energy Information Administration. (2020). 'Spot prices for crude oil and petroleum products [dataset]', https://www.eia.gov/dnav/pet/pet_pri_spt_s1_d.htm (last accessed: 7 April 2020).
European Commission. (2008). 'Restrictive measures', http://eeas.europa.eu/archives/docs/cfsp/sanctions/docs/index_en.pdf.
European Union. (2010). 'Council decision of 26 July 2010 concerning restrictive measures against Iran and repealing Common Position 2007/140/CFSP', *Official Journal of the European Union* (L195), pp. 39–73.
Eurostat. (2008). *Nace Rev. 2: Statistical Classification of Economic Activities in the European Community*, Luxembourg: Office for Official Publications of the European Communities.
Faccio, M. (2006). 'Politically connected firms', *American Economic Review*, vol. 96(1), pp. 369–86.
Fama, E.F. and French, K.R. (1993). 'Common risk factors in the returns on stocks and bonds', *Journal of Financial Economics*, vol. 33(1), pp. 3–56.
Fisman, R. (2001). 'Estimating the value of political connections', *American Economic Review*, vol. 91(4), pp. 1095–102.
GDELT Project. (2018). 'Global data on events, location and tone [dataset]', https://www.gdeltproject.org/ (last accessed: 1 November 2018).
Guidolin, M. and La Ferrara, E. (2007). 'Diamonds are forever, wars are not: Is conflict bad for private firms?', *American Economic Review*, vol. 97(5), pp. 1978–93.
Haidar, J.I. (2017). 'Sanctions and export deflection: Evidence from Iran', *Economic Policy*, vol. 32(90), pp. 319–55.
Heath, T. (2016). 'How a Trump presidency will affect 15 industries', *Washington Post*, 12 November.
Higgins, A. (2006). 'A feared force roils business in Iran: As hard-liners rise, shadowy Revolutionary Guard muscles in on airport and nabs energy deals', *Wall Street Journal*, 14 October.
Hsieh, C. and Moretti, E. (2006). 'Did Iraq cheat the United Nations? Underpricing, bribes and the oil for food program', *Quarterly Journal of Economics*, vol. 121(4), pp. 1211–48.
Hufbauer, G., Schott, J. and Elliott, K. (1990). *Economic Sanctions Revisited: Supplemental Case Histories*, Washington, DC: Institute for International Economics.
Iacus, S.M., King, G. and Porro, G. (2012). 'Causal inference without balance checking: Coarsened exact matching', *Political Analysis*, vol. 20(1), pp. 1–24.
Intrade. (2014). 'Daily price of contract 'US and/or Israel execute an overt airstrike against Iran by December 31, 2012' [dataset]', https://intrade.com/ (last accessed: 9 September 2014).
Jaeger, D.A. and Paserman, M.D. (2008). 'The cycle of violence? An empirical analysis of fatalities in the Palestinian-Israeli conflict', *American Economic Review*, vol. 98(3), pp. 1591–604.
Jahan-Parvar, M.R. and Mohammadi, H. (2013). 'Risk and return in the Tehran Stock Exchange', *Quarterly Review of Economics and Finance*, vol. 53(3), pp. 238–56.
JCPOA. (2015). 'Joint comprehensive plan of action', https://www.consilium.europa.eu/en/policies/sanctions/iran/jcpoa-restrictive-measures/ (last accessed: 1 August 2021).
Kaempfer, W.H. and Lowenberg, A.D. (1988). 'The theory of international sanctions: A public choice approach', *American Economic Review*, vol. 78(4), pp. 786–93.
Kerr, P.K. (2018). 'Iran's nuclear program: Status', Congressional Research Service Report RL34544, updated 19 September.
Knight, B. (2006). 'Are policy platforms capitalized into equity prices? Evidence from the Bush-Gore 2000 presidential election', *Journal of Public Economics*, vol. 90(4–5), pp. 751–73.
Loughran, T. and McDonald, B. (2011). 'When is a liability not a liability? Textual analysis, dictionaries, and 10-Ks', *Journal of Finance*, vol. 66(1), pp. 35–65.
MacKinlay, A.C. (1997). 'Event studies in economics and finance', *Journal of Economic Literature*, vol. 35(1), pp. 13–39.
Manacorda, M. and Tesei, A. (2020). 'Liberation technology: Mobile phones and political mobilization in Africa', *Econometrica*, vol. 88(2), pp. 533–67.
Marcial, G. (2016). 'Trump victory boosts demand for defense stocks', *Forbes*, 13 November.
Marinov, N. (2005). 'Do economic sanctions destabilize country leaders?', *American Journal of Political Science*, vol. 49(3), pp. 564–76.
Nasim Online. (2013). 'What is happening in third round of nuclear talks in Geneva', http://en.nasimonline.ir/archives/11665 (last accessed: 24 June 2014).

Neuenkirch, M. and Neumeier, F. (2016). 'The impact of US sanctions on poverty', *Journal of Development Economics*, vol. 121(1), pp. 110–9.

Organization of the Petroleum Exporting Countries. (2020). 'OPEC basket price [dataset]', https://www.opec.org/opec_web/en/data_graphs/40.htm (last accessed: 7 April 2020).

Pawlak, J. and Dahl, F. (2014). 'U.S. warns "time is short" as Iran nuclear talks make little progress', *Reuters*, 16 May.

Refinitiv. (2019). 'Datastream [dataset]', https://www.refinitiv.com/en/products/datastream-macroeconomic-analysis (last accessed: 14 February 2019).

Roodman, D., MacKinnon, J.G., Nielsen, M.Ø. and Webb, M.D. (2019). 'Fast and wild: Bootstrap inference in Stata using boottest', *Stata Journal*, vol. 19(1), pp. 4–60.

Schrodt, P.A. (2012). 'CAMEO conflict and mediation event observations event and actor codebook', http://data.gdeltproject.org/documentation/CAMEO.Manual.1.1b3.pdf.

Silver, N. (2016). 'Why FiveThirtyEight gave Trump a better chance than almost anyone else', *FiveThirtyEight*, 11 November.

Snowberg, E., Wolfers, J. and Zitzewitz, E. (2007). 'Partisan impacts on the economy: Evidence from prediction markets and close elections', *Quarterly Journal of Economics*, vol. 122(2), pp. 807–29.

Stecklow, S., Dehghanpisheh, B. and Torbati, Y. (2013). 'Assets of the Ayatollah: The economic empire behind Iran's Supreme Leader', *Reuters*, 11 November.

Stockholm International Peace Research Institute. (2016). 'SIPRI Arms Industry Database [dataset]', https://www.sipri.org/databases/armsindustry (last accessed: 21 June 2016).

Tehran Stock Exchange. (2011). 'Iran's capital market: How to invest in Tehran Stock Exchange (TSE)', http://www.tse.ir/cms/Default.aspx?tabid=82 (last accessed: 23 April 2014).

Tehran Stock Exchange. (2014). 'Market cap [dataset]', http://www.tse.ir/en/companylist.aspx (last accessed: 10 March 2014).

Tehran Stock Exchange. (2017). 'Shareholders [dataset]', http://www.tse.ir/market/shholder.aspx (last accessed: 3 January 2017).

Tehran Stock Exchange. (2018). 'Daily trading details [dataset]', https://tse.ir/en/archive.html (last accessed: 14 July 2021).

Tetlock, P.C. (2007). 'Giving content to investor sentiment: The role of media in the stock market', *Journal of Finance*, vol. 62(3), pp. 1139–68.

Torbati, Y. and Dehghanpisheh, B. (2015). 'Conglomerate controlled by Iran's Supreme Leader a winner from nuclear deal', *Reuters*, 21 October.

United Nations. (2010). 'United Nations Security Council resolution 1929', https://www.undocs.org/S/RES/1929%20(2010) (last accessed: 24 September 2017).

United States Treasury. (2013). 'Treasury targets assets of Iranian leadership', https://www.treasury.gov/press-centre/press-releases/Pages/jl1968.aspx (last accessed: 24 September 2017).

Warrick, J. and Rezaian, J. (2013a). 'Iran nuclear talks end on upbeat note', *Washington Post*, 27 February.

Warrick, J. and Rezaian, J. (2013b). 'In talks with Iran, reality tempers hopes on nuclear deal; optimism gives way to realism as diplomats ponder long odds of a reaching nuclear deal with Iran', *Washington Post*, 6 March.

World Bank. (2019). 'World Bank open data [dataset]', https://data.worldbank.org (last accessed: 17 September 2019).

Wroughton, L. (2013). '"Quite possible" Iran, powers can reach nuclear deal next week: U.S. official', *Reuters*, 15 November.