May 2018 No.372

On Target? The Incidence of Sanctions Across Listed Firms in Iran

Mirko Draca, Jason Garred, Leanne Stickland
and Nele Warrinnier

WORKING PAPER SERIES

Centre for Competitive Advantage in the Global Economy

Department of Economics

On Target? The Incidence of Sanctions Across Listed Firms in Iran

Mirko Draca* Jason Garred† Leanne Stickland‡ Nele Warrinnier§

July 21, 2017

Abstract

A central premise of current international sanctions policy is targeting, that is, concentrating the impact of sanctions on specific, politically influential groups in the sanctioned country. However, many economic factors make it difficult for senders of sanctions to hit these targets. We offer evidence on the efficacy of targeting in the case of Iran, where sanctions aimed to affect a well-defined set of political entities through their economic interests. Our identification strategy focuses on the process of negotiations for sanctions removal. We find that stock returns of firms owned by targeted political groups and firms unrelated to these groups both react positively to information indicating progress in diplomatic negotiations. However, these effects are significantly larger for firms owned by targeted groups. This evidence suggests that good news about sanctions relief yielded particularly large economic benefits for targeted political entities, consistent with the 'income targeting' goal of sanctions policy against Iran.

We thank Abel Brodeur and seminar participants at Paris School of Economics, LSE, Helsinki, Carleton, King's College London, Ottawa, 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 and the LICOS Centre for Institutions and Economic Performance 25 Years Workshop. We also thank Femi Amoo, Terry Cheng, Divya Deepthi and Neel Ocean for excellent research assistance.

*Department of Economics, University of Warwick

†Department of Economics, University of Ottawa

‡National Audit Office, UK Government

§Faculty of Economics and Business, KU Leuven

1 Introduction

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 program for the development of nuclear technology. Other prominent recent examples of sanctions have included measures leveled 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 construction of modern sanctions regimes is that they should be targeted as much as possible to the economic interests of elite decisionmakers.[1]

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. For example, the UN's Oil-for-Food program 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). 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.

In this paper, we consider the case of Iran, which was subject to multilateral sanctions related to its nuclear program until 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. At the same time, the effect of sanctions on the rest of the Iranian population would have been minimized, and so sanctions removal should have had little impact on the welfare of these other actors.

Our study attempts to understand the effectiveness of targeting of sanctions in Iran through a stock market event study covering this period of diplomatic progress. We use information on owners of firms listed on the Tehran Stock Exchange (TSE) to assemble two

[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."

firm portfolios. The 'target' portfolio includes assets of conglomerates controlled by two political actors with strategic roles in decisions about Iran's nuclear program: the Supreme Leader and the Islamic Revolutionary Guard Corps (IRGC). Importantly, the targeting of these interests within Iran was an explicit goal of sanctions policy by the US and its partners in the sanctions regime. We also construct a 'non-target' portfolio consisting of firms in which these two actors had no known ownership stake.

Our empirical strategy has two complementary elements designed to capture market responses to progress in diplomatic negotiations about sanctions relief. We first explore a compelling case study: the stock market reaction to an unexpected 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 this particular round of negotiations accruing over the TSE's weekend break. As such, we are able to cleanly capture the TSE market response to the most important turning point in the negotiations for sanctions removal.

The second part of our empirical strategy is then based on a high-frequency, news-based measure derived from the comprehensive Global Database of Events, Language and Tone (GDELT). The GDELT database provides 'near real-time' political event data produced using automated reviews of coverage from 37,000 sources. The key feature of GDELT for our analysis is that it is able to consistently identify events specifically related to negotiations between Iran and its sanctioners as natural language objects across news articles, allowing us to track the development of sanctions negotiations on a daily basis. We use coverage measures derived from GDELT data to evaluate stock market responses to daily variation in news about diplomatic negotiations over three years between the revival of multilateral negotiations in April 2012 and the conclusion of a final agreement in July 2015. Importantly, because the news that we consider relates to the lifting of the sanctions regime as a whole, both of our approaches (the November 2013 case study and the analysis using GDELT) allow us to relate our estimated effects to the probability of comprehensive sanctions relief rather than a single sanctions instrument.

Our results show a consistent pattern of responses by firms in the two portfolios to our posited information shocks. Specifically, in both of our empirical exercises, we find that 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, these effects are significantly larger for listed firms owned by targeted political entities.

For each day of a two-day event window, the estimated positive effect of the Geneva deal on the returns of non-target firms is approximately 30% of the size of the standard deviation of returns across all firms and days in the sample. For firms in the target portfolio, the estimated abnormal returns from Geneva exceed this by an additional 65% of a standard deviation, for a total effect of approximately two percentage points per day. Meanwhile, a one standard deviation jump in our GDELT daily news coverage measure is associated with a 0.06 percentage point increase in returns for non-target firms, with an additional increase of approximately the same magnitude accruing for the target group. Importantly,

this relationship between sanctions news and returns holds strongly even outside the period immediately around the November 2013 Geneva deal.

This finding of systematic comovement between target firm returns and events relating to sanctions removal is notable for providing evidence that sanctions against Iran succeeded in inducing some degree of 'income targeting'. Even though we cannot say whether decisionmakers were driven by this incentive when negotiating the removal of sanctions, our empirical evidence suggests that such an incentive existed.

A potential identification challenge for our research design is that news about sanctions might proxy for changes in the likelihood of conflict between Iran and the West, which might also differentially affect target and non-target firms. We therefore perform two additional exercises in order to assess whether our news-based measures include information about changes in war risk, and how firms in our sample respond to the risk of conflict. We first show that while both target and non-target firms are sensitive to a political betting market measure of Iran's probability of direct military conflict with the US or Israel, the target portfolio is not differentially responsive to this information. We then use a portfolio of arms industry firms listed on stock markets outside Iran to investigate whether there is a strong overlap between sanctions news and the changing risk of war. The returns of these war-sensitive firms show no significant co-movement with sanctions news, suggesting that news about negotiations is well-separated from updates to expectations of conflict risk.

We draw on and contribute to a number of literatures. Firstly, we make 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, Schott and Elliott (1990) and focuses on whether the offending policies of the receiver are actually altered after sanctions are imposed. Very few papers zoom in to consider the incidence of sanctions within the receiver, and our microeconomic assessment of the 'income targeting' goal of sanctions is particularly innovative. Perhaps the closest study to ours is Ahn and Ludema (forthcoming), which concludes that recent US and EU 'smart sanctions' on specific Russian companies have had substantial negative effects on those firms' financial performance. Haidar (2017) also studies sanctions on Iran, and finds heterogeneous effects of these sanctions on exporting firms depending on characteristics such as exporter size. Marinov (2005) finds that leaders of countries subject to economic sanctions are more likely to lose power. This is in line with Allen (2008), which notes a positive relationship between sanctions and antigovernment activity. Neuenkirch and Neumeier (2016) investigate the effect of US sanctions on the poverty gap in sanctioned countries (a measure of the incidence and severity of poverty), and argue that sanctions lead to an increase in poverty.

Secondly, this paper is one of a number of recent studies using variation in stock market returns to study topics in political economy. In particular, these studies focus on questions such as private benefits of civil conflict (Guidolin and La Ferrara (2007) on diamond mining in Angola), the economic implications of covert foreign intervention (Dube, Kaplan and Naidu (2011) on post-war CIA operations), the detection of illegal arms trade (Della Vigna

and La Ferrara 2010), and the value of political connections (Fisman 2001, Faccio 2006, Braggion and Moore 2013, Coulomb and Sangnier 2014 and Acemoglu et al. 2016).

Finally, we also contribute to a growing literature on the political economy of the Middle East. In a similar vein to our study, the research design of Durante and Zhuravskaya (forthcoming) uses the timing of news events to study Israeli military strategy. Acemoglu, Hassan and Tahoun (forthcoming) argue that more intense street protests in Egypt restricted access to excess rents for firms connected to political incumbents. Jaeger and Paserman (2008, 2009, 2015) and Jaeger et al. (2012) examine the dynamics of violence in the Israeli-Palestinian conflict, while Berman, Shapiro and Felter (2011) provide insight into Iraq's insurgency.

The paper is structured as follows. Section 2 presents a conceptual framework to motivate our study and its research design. Section 3 provides background information on the history of sanctions on Iran and the political entities targeted by these sanctions. Section 4 describes the data we use, and Section 5 then presents our empirical analysis. Section 6 assesses the possible alternative interpretation of our results in terms of changes in the risk of conflict, and Section 7 offers concluding remarks.

2 Conceptual framework

In this section, we first set out a simple static framework to clarify why countries might want to impose targeted sanctions, and why the effectiveness of targeting is an important economic question. As part of this, we frame the effects of sanctions as creating an incentive ('carrot') for policy changes in a sanctions-receiving country, distinguishing between targeted and non-targeted sub-groups within the receiver. We then introduce 'blunt instruments' in the realized 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 sub-groups.

We then 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. In particular, the spread of impacts suggested by the model lets us tie our sanctions bluntness parameter β to the results of a difference-in-difference model that compares the asset values of targeted and non-targeted groups within the receiver over time.

2.1 Setup

We consider a scenario in which sanctions are imposed in order to induce changes in a particular policy in the sanctioned country, since this is the usual publicly stated goal of multilateral sanctions on Iran. 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 the group p within R benefits from sanctions relief, the removal of sanctions is effectively a reward for ending the policy that harms S.[2] Say that the cost of sanctions in R is y, and that this cost is spread additively across group p and the group o, constituting the rest of the population of R, so that $y_p + y_o = y$. 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; Y_p, the income of group p, and $Y^R = Y_p + Y_o$, the aggregate income of R. The relevance of Y_o to the welfare of p might be due to altruism or 'good governance' by group p, or on the need for support from the population as a whole in order to sustain the political regime. If p places a weight $\alpha - 1$ on Y_p, where α may differ from 1, then welfare of p is equal to $b_p + \alpha Y_p + Y_o$.

To understand the potential attractiveness of targeting to S, assume that the cost of sanctions to S rises in y, so that increasing the cost of sanctions in the receiving country also increases the cost to the sender.[3] Then consider the scenario in which S has complete control over their incidence across o and p. In this case, S solves the problem:

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

As long as $\alpha > 1$, the solution to this problem is to set $y_p = \frac{1}{\alpha} b_p$, $y_o = 0$. 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.[4]

2.2 'Blunt instruments'

Now imagine that S cannot control the incidence of sanctions within R, perhaps because of the bluntness of the available sanctions instruments, or the asymmetric capacity for sanctions avoidance across groups within R. In particular, say that an attempt at a 'fully targeted' regime of cost y actually leads to incidence $y_p = \beta y$, $y_o = (1 - \beta) y$ for some $\beta \in [0, 1]$ that is exogenous to S. 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 the income costs of sanctions, while $\beta = 1$ implies that the full income costs fall on the political elite only.

In this 'bluntness' case, the chosen sanctions regime will be weakly costlier than in the scenario where S can control sanctions incidence. In particular, the solution for S is now to set $y = \frac{1}{\alpha\beta + (1-\beta)} b_p$, which is a weakly greater total cost than under the targeted regime,

[2] Note that because of the empirical context we study (in which sanctions are in place and the sender and receiver are already negotiating sanctions relief), we 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.

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

[4] Note that in this static framework, we implicitly assume 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 incidence of costs that arises when sanctions first hit is different from the incidence of benefits when sanctions are lifted.

where $y = \frac{1}{\alpha} b_p$. 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.

The above framework may be extended to accommodate other possible scenarios, without substantially changing its main implications. For example, it might be the case that there are two groups within the receiver with potential influence over the policy, and depending on its incidence across these two groups, sanctions can change the balance of power between them, as in Kaempfer and Lowenberg (1988). If the relative strength of the group with a stronger preference for the policy is eroded by the relative magnitude of its economic losses from sanctions, then it will again be optimal for the sender to target that group, and the magnitude of β (where β now determines the relative incidence of sanctions across these two policymaking groups) will again be of economic interest. A similar argument applies if the goal of sanctions is to change the regime in power rather than to affect an individual policy: the sender will wish the sanctions to cause greater economic harm to groups in the government relative to the opposition, as long as economic power is a source of political power within R.

2.3 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.[5] 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 isolates the effects of these negotiations on the costs of sanctions to the targeted political elite p and to the non-elite o. We use stock returns 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 comove positively with news of progress in negotiations, and these asset values are unrelated to any benefits from the policy itself, then we interpret this comovement as evidence that $\beta \neq 0$. Practically, this implies that the group p benefits from expected sanctions relief such that it has not escaped bearing an income cost from the sanctions. If the values of assets of the non-elite group o move together with these shocks, we similarly consider this to be evidence that $\beta \neq 1$; that is, the elite group do not bear

[5]Note here that we do not consider the more complex situation in which the bluntness parameter β itself evolves over time.

the full costs and there is some sharing of costs between p and o. Finally, we derive further information about β by comparing the relative magnitudes of these changes in the values of the two groups' assets. As we outline in Section 5, this maps into our empirical difference-in-difference model such that we are able to provide a broad empirical assessment of the success of the realized sanctions policy in terms of the posited targeting bluntness parameter β.

3 Background

3.1 Sanctions on Iran and political negotiations

While the United States has maintained economic sanctions on Iran since soon after its 1979 revolution,[6] robust multilateral sanctions prompted by Iran's nuclear program were imposed only from the mid-2000s. 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 whether Iran's nuclear program had no military dimension. The resolution called for Iran to suspend enrichment and reprocessing activities associated with its nuclear program. Sanctions were then imposed in December 2006 and tightened in two subsequent resolutions in 2007-08, 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 by the US, EU and other actors constituted both 'smart sanctions' and more comprehensive measures. For instance, the Security Council resolutions asked states to freeze the assets of specific firms, individuals and other organizations directly involved in Iran's nuclear program, and also to "exercise vigilance" over the foreign activities of Iranian financial institutions. Sanctions were publicly portrayed as targeting the economic interests of actors with political influence inside Iran.[7] However, some observers argued that sanctions were ineffective, or even strengthened defenders of the nuclear program within Iran: for example, the 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."[8]

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

[6] US sanctions on Iran have included a variety of measures relating not only to weapons proliferation, but also to terrorism and human rights.

[7] For example, after the US cut off 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." (Wall Street Journal 2006)

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

Council (China, France, Russia, the UK and the US) as well as Germany. This process failed to reach agreement on sanctions relief in exchange for policy change 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 US and EU. These new sanctions imposed severe limitations on Iran's international financial access; for example, the Belgium-based organization 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 program resumed with a meeting in Istanbul in April 2012, and were characterized 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 program. 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. Finally, 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.

3.2 Targeted political entities

From an early stage in the development of the multilateral sanctions on Iran, the sanctions were focused heavily on entities (i.e. individuals, companies or units of government) that were perceived to either provide direct infrastructural support to the nuclear programme or have a major political influence over the direction of Iran's nuclear strategy. Specifically, the key entities with special political influence over the programme were identified as the Islamic Revolutionary Guard Corps (IRGC) and the Supreme Leader, who controls a conglomerate of companies known under the name of EIKO (the Execution of Imam Khomeini's Order) or 'Setad'. We discuss the nature of each of these actors' influence over the nuclear programme and the structure of their economic interests in turn.

3.2.1 Islamic Revolutionary Guard Corps (IRGC)

The IRGC is a branch of the Iranian military with the explicit political role of guarding the post-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.

One channel of the IRGC's influence over the nuclear programme is its role in the technical administration of weapons development. It runs Iran's main domestic weapons supplier, the Defense Industry Organization, and was in direct control of many of the scientific units and physical assets that underpinned the nuclear programme. A prominent example of this is the military site of Parchin, which was thought to house key nuclear weapons infrastructure, but was barred from full IAEA inspection from 2005 onwards. The fact that the military held direct decision-making power over the right to inspection for sites such as Parchin was acknowledged by the head of Iran's atomic energy organization, Fereydoun Abbasi-Davani, in 2012 (Reuters 2012).[9]

A further important example of an IRGC-controlled nuclear asset was the Fordow fuel enrichment plant, constructed as an underground tunnel system on the grounds of an IRGC military base. This facility became a focus of nuclear diplomacy after its disclosure in September 2009, with a significant wave of sanctions covering 15 IRGC-controlled firms following in June 2010 as part of UN Security Council Resolution 1929. The conversion of Fordow into a 'technology center' stripped of many nuclear functions was an notable part of the final multilateral agreement to lift sanctions (Katzman and Kerr 2017).

The income base of the IRGC provided many natural targets for sanctions. The IRGC has an extensive economic presence across sectors in Iran, derived both from its participation in military industries and its financial wing providing services for current personnel and veterans. Sanctions were explicitly motivated as affecting these interests, with the US Treasury citing the economic roles of key IRGC-controlled companies in its sanctions designations.

3.2.2 Supreme Leader and Setad

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 37 firms in a network associated with the organization 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 Iran's head of state, Supreme Leader Ali Khamenei, and noted that Setad's revenues allowed for Khamenei's financial independence from Iran's parliament and national budget process (Reuters 2013).

Reports 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

[9]Hassan Rouhani also wrote of the contentious and lengthy internal debate about inspections in his detailed 2011 memoir of Iranian nuclear negotiations. See Kerr (2017) for notes on Rouhani's 1,200-page Persian-language memoir "National Security and Nuclear Diplomacy". Rouhani wrote that "inspections of military centers such as Parchin [were] debated for many months inside the country.... There was serious opposition to the Agency's request to inspect Parchin; the nation's domestic political climate was vigorously opposed to inspectors inspecting Parchin and military centers in general."

Reuters (2015) indicated that the organization was targeted because the US "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 the time of Setad's designation, nuclear talks were deadlocked and action was 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)". (Reuters 2015) These reports also suggested that the lifting of sanctions against Setad yielded tangible economic benefits for the conglomerate, such as the facilitation of international business deals.

4 Data

4.1 Stock returns data

We collect information on stock returns of listed Iranian firms by web scraping daily data from the web site 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 quartile of financial access (share of companies outside the top 10 in market capitalization), the third quartile of financial depth (market capitalization as a share of GDP) and the second quartile of financial efficiency (shares traded as a proportion of total market capitalization), relative to other countries.[10] The TSE trades for three hours per day (9:00 to noon), five days per week (Saturday through Wednesday).

Our full 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 breakdown or hiatus in talks) and because the majority of sanctions were in place by early 2012. This latter point facilitates an event study research design centred on the effects of comprehensively lifting sanctions. We expand on these points further in Section 5.2.

As of the beginning of our sample period, there were 325 firms listed on the TSE. As discussed in the next subsection, our sample is made up of two firm portfolios, which together include 137 of these 325 firms. Daily returns for each firm are calculated by subtracting a stock's closing price for the previous day from its closing price for the current day and dividing by the closing price for the previous day. We multiply daily returns by 100 so that

[10]The TSE ranks very similarly to Egypt by these measures, based on data from the Egyptian Exchange. Studies of market efficiency (Jahan-Parvar and Mohammadi 2013) indicate 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.

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.

To classify firms into industries, we use a concordance of firms to sectors in an English-language guide to investment in Iran published by the TSE itself (TSE 2011), supplemented in several cases by hand-coding for sample firms missing from this list. We then concord the sectors defined by the TSE guide to 2-digit NACE codes, amalgamating some of these codes into more broadly defined industries so that there are no sectors in which only one of the two portfolios is represented. The resulting 15 industries are listed in Table 1.

4.2 Ownership data

4.2.1 Target portfolio

In line with the discussion in Section 2.3, 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 behind these conglomerates gained from sanctions relief. We also wish to exclude firms with business directly related to Iran's nuclear program. This is so that we do not confound the expected effects of sanctions relief with the expected impact of changes in Iran's nuclear policy itself, again as discussed in Section 2.3. To identify a suitable portfolio of target firms, we rely on two main sources of data: information in sanctions-related documents from the US, EU and UN, and information on shareholders of firms listed on the TSE, translated from the Persian-language TSE web site.

We first exploit the fact that when 'smart sanctions' (such as an asset freeze) are applied to a particular firm, the reasons for doing so are normally documented by the sanctioner. We therefore survey all US Treasury press releases announcing that US sanctions of this type have been extended to one or more additional companies. These sanctions are often directly motivated by the fact that the firms concerned are either part-owned or fully owned by IRGC or Setad. When this is the case for a given firm, this is noted in the US Treasury press releases.

Because both of these organizations control large conglomerates of companies, their ownership structure is complex; see Figure 1 for a diagram, published by the US Treasury, depicting several layers of firms controlled by Setad. Each of these conglomerates is comprised of a mixture of stock-market listed and private, unlisted companies. Thus, although information from the US Treasury allows us to identify 69 firms and other entities partly or fully owned by the IRGC and/or Setad, only six of these are listed on the TSE as of the beginning of the sample. However, using TSE data on the shareholders of each listed firm,[11] we identify 47 additional firms for which at least one of the 69 firms identified by the US

[11]Shareholder data is from April 14, 2012, the first day of our sample period. Changes in ownership over time might mean that later in the sample period, the returns of this portfolio may not fully reflect actual gains and losses for current IRGC and Setad assets. However, our results are robust to instead using shareholder data from the final trading day in the week before the Geneva deal (November 20, 2013), which falls in the middle of the full sample period.

Treasury is a shareholder.

Documents identifying firms subject to 'smart sanctions' are again of use in identifying firms directly associated with the Iranian nuclear program. Specifically, we drop firms from the target portfolio if a document from the US, EU or UN notes that the firm conducts business directly relating to Iran's nuclear program, such as involvement in reactor construction. This excludes three firms from the 53 identified above, leaving a portfolio of 50 firms in which the IRGC and/or Setad have ownership stakes.

4.2.2 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 272 listed firms that were not identified by the procedure documented in the previous subsection. However, while this excludes the 53 firms with the most direct 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.

First, we again take our list of 47 firms identified as IRGC or Setad assets by the TSE shareholder data, and search for additional firms on the TSE that have one or more of these 47 firms as shareholders. In other words, we identify a set of firms that are two layers of ownership removed from the firms identified by the US Treasury. We then continue this process, identifying additional ownership layers using the TSE data, until we have identified all listed firms that are two or more ownership layers below the firms in the US Treasury data. We exclude these 147 firms from the non-target portfolio.

Second, if a firm in our sample has been identified as an IRGC or Setad asset by unofficial sources – a 2013 Reuters investigative report on Setad (Reuters 2013) or a 2010 American Enterprise Institute report on IRGC's involvement in the Iranian economy (Alfoneh 2010) – but is not already in our target portfolio, then we drop it from the non-target portfolio. We also browse the web sites of two large investment companies in our target group (Bahman Group and Ghadir Investment) and drop listed firms that are subsidiaries of either of these companies but do not appear in the target portfolio. We then again use the TSE shareholder data to construct layers of ownership below the firms identified by these additional sources, and remove all firms within this ownership structure from our non-target portfolio. This procedure eliminates another 31 firms, leaving 94 non-target firms. Finally, we again drop firms directly associated with the Iranian nuclear program, according to documents from the US, EU and UN, which excludes seven additional firms. The final non-target portfolio is made up of 87 firms.

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 capitalization: between one fifth and one quarter in each case.[12] The first two moments of the daily return variable by firm-day are also very similar across the two portfolios. Accounting data from Orbis for 2012, which is available for only a subset of firms,

[12]Reliable archival data on TSE market capitalization by firm is unavailable, so we instead collected current market capitalization data by hand in March 2014.

suggests 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. We account for these differences in sectoral composition in some of our empirical specifications below.

4.3 Event data

4.3.1 GDELT overview

We use measures derived from the Global Data on Events, Location and Tone (GDELT) database as proxies for day-to-day progress in political negotiations during the sample period. 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 reports in the world's news media. The GDELT Project extracts information from the text of news articles covering more than 37,000 online news sources in 100 different languages, automatically updated every 15 minutes. The key innovation of GDELT is that it stores this information as consistently coded event data, where the events are automatically defined by NLP algorithms. The scale of the global media coverage achieved by the GDELT Project yields a database of over a quarter-billion distinct political events that have occurred since 1979.

Articles collected by GDELT are first machine-translated into English, after which NLP algorithms extract information about events, actors, locations and tone and classify them according to the Conflict and Mediation Event Observations (CAMEO) taxonomy.[13] The codified metadata is then stored in the GDELT Project archives and released as an open, 'near real-time' data stream.

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). Actions are classified into CAMEO's hierarchical event typology at the two-digit, three-digit and four-digit levels, covering all political acts of verbal and material mediation and conflict. A list of political event classes defined at the two-digit level is provided in Appendix Table A1.

4.3.2 Examples of GDELT automatic event coding

In order to give more insight into the machine-learning classification algorithm employed by GDELT, consider the following examples from the CAMEO manual (Schrodt 2012):

[13] GDELT uses the Textual Analysis by Augmented Replacement (TABARI) software for extracting events from news articles. More specifically, TABARI parses every sentence of news articles into a subject, verb phrase and a direct object. Both subject and object are than checked against all actor codes stored in the CAMEO taxonomy and stored into a string of actor codes if a match occurs. Similarly, the verb phrase is checked against a dictionary of verb phrases and stored as a CAMEO event code if a match occurs. Locations of the actors are defined by GDELT as the geographical location mentioned in the article that is fewest words away from the actor.

> Example 1: "The **Russian Foreign Minister Sergei Lavrov** *said* here Saturday that he *urges* Iran and the EU trio (France, Germany, and Britain) to *reach* an *agreement* in their talks on Iran's nuclear program."

Based on the verb phrases identified in the sentence (in italics), GDELT will determine that this event is an appeal by a certain actor to settle a dispute between two other actors, and hence will classify it under the three-digit 'Appeal to others to settle dispute' category. The subject actor is identified as the one who initiates the action (in bold), in this case Sergei Lavrov of Russia. The object actor is identified as the one who receives the action (underlined) and in this case is five different entities (Iran, the EU, France, Germany and Britain). Given the presence of five different object actors, GDELT will enter five separate events into its database with five unique event identifiers.[14]

> Example 2: "The IAEA has *dispatched inspectors* to Esfahan Uranium Conversion Facilities (UCF) in central **Iran** to monitor resumption of peaceful nuclear work at the plant."

Based on the verb phrase "dispatched inspectors" (in italics), GDELT will determine that two distinct events took place. Firstly, there is an actor (underlined) who sends inspectors to another actor (in bold), which will be classified under three-digit CAMEO code 090 'Investigate, not specified below'. Secondly, there is an actor (in bold) who receives inspectors from another actor (underlined), which will be classified under four-digit CAMEO code 0862 "Receive inspectors".

Any event recorded with exactly the same date, subject actor code, object actor code and event code is treated as a single event and given the same unique event identifier in the GDELT database. Hence, GDELT treats events as meta-objects across time and news articles. In other words, a certain event can be reported on in multiple articles and over multiple days, but this information is linked by GDELT to the original event at the time it occurred. To keep track of how much coverage a certain event receives, GDELT reports a count measure of all source articles mentioning the event.

4.3.3 Aggregating GDELT events

For our analysis, we generate a time-series measure capturing the diplomatic negotiations between Iran and the P5+1, by extracting 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.[15] For our baseline measure, we only include events in two-digit CAMEO categories into which diplomatic negotiation events are most likely to have been classified: "03 - Express intent to cooperate"; "04

[14]Note that this coding rule probably increase the importance of multilateral diplomatic negotiations in our news coverage measure relative to bilateral events involving Iran and one other actor. Since these bilateral events seem less likely to be directly related to sanctions negotiations, this may actually decrease measurement error in our coverage variable.

[15]We use GDELT data current as of September 2016.

- Consult" and "05 - Engage in diplomatic negotiations".[16] Because our goal is to measure investors' reactions to these events on the TSE, we assign all events occurring over the TSE's weekend (Thursday and Friday) or during holidays to the following trading day.

We measure events on an intensive margin, using GDELT's count measure of source articles that mention a certain event, which can be interpreted as proxy for the event's importance. We sum the total number of article mentions of all events classified to a certain day in our three two-digit CAMEO categories, and standardize this measure across all of the days in our sample. Before standardizing, 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.

Figure 2 displays the variation in this measure over the course of the sample period. The figure suggests that there are many spikes in this variable; our measure thus appears to identify a number of days over the course of the sample period with particularly important events involving Iran and P5+1 countries. An examination of the top fifteen days by this measure (Appendix Table A2) suggest that these events are indeed relevant to negotiations about removal of Iran sanctions. Several of these events are labelled in Figure 2.

4.4 Other data

In order to test for possible war-related explanations for our results, 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.[17] 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. Hence, the price of the contract yields a market-based estimate of the probability of an airstrike, which we use as a proxy for the odds of military conflict. Figure 3 shows the evolution of this contract's daily price during 2012.[18]

Second, we use the Datastream database to extract stock price data for a set of firms for which profit expectations are likely to be sensitive to the odds of war in the Middle East. In particular, we gather this information for all publicly traded firms in the Stockholm International Peace Research Institute (SIPRI) 2012 Arms Industry Database. The SIPRI database identifies the hundred largest arms-producing and military services companies in

[16] We judged the neighbouring categories "02 - Appeal" and "06 - Engage in material cooperation" to be insufficiently related to diplomatic negotiations for inclusion in our baseline measure. However, as we show in Table 5, our results are robust to including events in all CAMEO categories.

[17] 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 1 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.

[18] We do not display price data after December 22, 2012, since US trading on Intrade was suspended as of December 23.

2012, but since some of these firms are privately owned, we observe stock price data for 66 of the companies on this list as of the beginning of our sample period. We use daily stock price information from Datastream to calculate daily returns for these firms over the course of the period of study.

5 Empirical analysis

5.1 Case study: Geneva deal

To more clearly introduce our event study 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 delivers some useful advantages for our analysis. The final round of negotiations for the interim agreement began at 17:00 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. During this time, encouraging reports emerged about the progress of talks, which continued in Geneva over both Thursday and Friday.[19] When the TSE re-opened at 9:00 on Saturday November 23, there was thus a large amount of information on which the market could act. The final agreement was then officially announced at 5:30 Tehran time on Sunday November 24, in advance of the TSE opening again at 9:00 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 follows:

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

where i indexes firms, j represents industries and the time index t is at the day level. Here, R 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. This model can be interpreted as a variant of the constant mean return model in the event study literature (MacKinlay 1997) because of the inclusion of firm fixed effects. This is because we are effectively demeaning returns on a firm-specific basis, analogous to an abnormal return.

Our earlier conceptual framework implies the existence of three scenarios for the impact of targeted sanctions. The first is a case in which the sender can fully control the incidence of sanctions as felt by the receiver country; i.e. a situation in which the parameter β in our model equals one. This results in a 'minimum-cost sanctions policy' through which sanctions

[19] These included several Twitter updates from the Iranian foreign minister, with text including "Positive atmosphere", "Talking details", "Room for optimism" and a notice that the parties had reached agreement on the issue of nuclear enrichment (Nasim Online 2013).

affect the income of the targeted group exclusively. In the second scenario, $\beta = 0$, so that sanctions have no effect on the income of targeted political entities. In the third scenario, in which β lies between zero and one, there is an impact of sanctions on both groups, with the relative magnitude of this impact depending on the magnitude of β.

We argue that the good news about progress towards a sanctions deal in Geneva should imply higher future income 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. Therefore, in the first scenario where $\beta = 1$, and only the targeted political elites bear the costs of sanctions, we should observe that in the difference-in-difference model $\phi = 0$, since firms that are not owned by targeted political groups should not be affected by sanctions. In the second scenario where $\beta = 0$, we should observe that $\phi + \theta = 0$, because assets of IRGC and Setad should not change in value due to progress in negotiations.

The rejection of both of these hypotheses, however, would provide evidence in favour of the third scenario, in which both the targeted 'elite' group as well as non-elite interests in Iran benefit from sanctions relief. This would imply that the costs of sanctions were shared as a result of the 'blunt instruments' problem. The value of θ, the additional effect of Geneva on firms that are assets of targeted political actors, would then provide suggestive evidence of the magnitude of the bluntness parameter β. Of course, there is an important caveat to these interpretations: that we observe returns to Geneva only for listed firms. If other assets of IRGC, Setad or the owners of non-target firms are affected by sanctions relief in different ways from the stock price movements we observe here, then the net effect on these actors might be of a different magnitude or even a different sign from the one we observe. Alternatively, taking account of the impact of sanctions on Iranians who do not own stakes in non-target firms might also substantially affect our conclusions.

With these limitations in mind, we present the results of estimating equation (1) for a sample period that includes the two days November 23 and 24, 2013 and the sixty trading days 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.[20] As shown in column (1) of Table 2, we find that for non-target firms, the Geneva deal was associated with a return that was 0.610 percentage points above the mean return for those firms, with an additional effect of 1.294 percentage points for target firms. Both of these estimates are statistically significant when standard errors are clustered by firm.[21]

In column (2) of Table 2, we 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) displays the results of regressions in which dummies and interactions for November 23 ('day 1') and November 24 ('day 2') are included separately. For both

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

[21] Using the much stricter strategy of clustering by the 15 industries in our sample, and using the wild bootstrap approach of Cameron, Gelbach and Miller (2008) due to the small number of clusters, produces p-values of 0.072 for our estimated difference-in-difference coefficient and 0.196 for the coefficient on the Geneva dummy.

portfolios of firms, the estimated impact is similar across the two days: for non-target firms, this is 0.521 on day 1 and 0.700 on day 2, while our difference-in-difference estimates are 1.442 for day 1 and 1.147 for day 2.

Because November 23 and 24 are at the beginning of the trading week, when stock returns might differ systematically from other days, we rerun our baseline regression substituting firm-day-of-week fixed effects for firm fixed effects, and find very similar results (see column (3)). Another potential issue with our baseline specification is that firms in the target and non-target groups are somewhat unevenly distributed across industries, as shown in Table 1. Along with the promise of potential changes in the future evolution of sanctions, the Geneva deal also suspended some industry-specific sanctions (on crude oil, petrochemicals, precious metals, the automotive industry and precious metals). Our results might thus be capturing an asymmetric industry-level effect of the Geneva deal itself rather than an effect of the multilateral sanctions regime as a whole. However, in Table 2 column (4), we add industry-Geneva interactions to our baseline specification in order to control for industry-specific effects of Geneva, and find that our estimate of the differential effect of Geneva on firms in the target portfolio is almost unchanged.

Based on our discussion above, these results suggest that both target and non-target firms benefit from sanctions relief, but that the impact on target firms is significantly larger. The estimated average effect of each day of the Geneva event on non-target firms is approximately 30% of the size of the standard deviation of returns across all firms and days in the sample. However, the estimated abnormal returns for target firms are approximately 65% of a standard deviation higher than this. This is evidence in favour of the possibility that targeting of sanctions was neither a complete success nor a complete failure, but that firms owned by targeted political entities were more affected by sanctions than other listed firms.

The statistical significance of this difference between the returns of target and non-target firms suggests that such a difference in returns between these groups 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 overestimate the probability that we would observe such a large difference even if $\theta = 0$. 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 coefficients on the interaction term of each of these 388 regressions in Figure 4, highlighting the 'true' Geneva result in red. It is apparent 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 leading up to 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.

Finally, we 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 Geneva dummy, 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 * one-day dummy interactions, along with their 95% confidence intervals, in Figure 5.

Throughout the two-week period starting November 23 (to the right of the vertical line in 5), the difference-in-difference coefficient for each day is either positive (though not always statistically significant) or very close to zero. 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 temporary surge of pessimism about the potential future incidence of sanctions during the pre-deal negotiations period. 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.

5.2 GDELT news coverage analysis

5.2.1 Basic news coverage model

For our main empirical exercise, we extend our sample period to the full range of dates discussed in Section 4.1: April 14, 2012 to July 15, 2015.[22] 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 subperiods (quarters). Second, significant changes were made to the multilateral sanctions regime over the several months before April 2012. As discussed in our introduction, an identification strategy based on progress towards comprehensive sanctions relief rather than the imposition or removal of particular instruments allows 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 use a daily event measure sourced from the GDELT event dataset. Our baseline measure of progress is a standardized variable based on the number of news articles about relevant events, as described in Section 4.3. We amend the model of equation (1) by replacing our

[22] As discussed in Section 4.3, we exclude February to August 2013 because of a technical issue with GDELT article collection during this period.

Geneva dummy with this continuous news coverage variable as follows:

$$R_{ijt} = \alpha_i + \phi Coverage_t + \theta Target_i * Coverage_t + u_{ijt} \quad (2)$$

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 3 displays the results of estimating this specification. The estimates imply that an increase in our news coverage measure by one standard deviation is associated with a return that was 0.068 percentage points above the mean return for non-target firms, and an additional 0.066 percentage points above the mean for target firms. Both of these estimates are statistically significant at the 1% level.[23]

It is possible that our news coverage measure rises over the course of the sample for reasons other than progress in negotiations (such as better measurement in the GDELT data), which could bias our estimates upward if average stock returns also rise. Therefore, in column (2) we use within-quarter variation (i.e. a comparison of abnormal returns as measured within a particular quarter and year, such as January to March 2014) to arrive at our estimates. To do this, we substitute firm-quarter fixed effects for the firm fixed effects in equation (2).[24] This has little effect on the estimated coefficients and standard errors.

In column (3), we further refine our fixed effects to 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. This is our preferred specification, and we continue to use firm-quarter-day-of-week fixed effects in all subsequent specifications. However, our estimates change little from those of column (2). Moreover, in all of our specifications so far, both estimated coefficients of interest are positive and statistically significant. As with our Geneva results, this lends support to the hypothesis that targeting of sanctions on Iran was not entirely successful, but that sanctions' impact on firms owned by targeted actors was larger than their effect on other listed companies.

Finally, in column (4) of Table 3, we add interactions of our coverage measure with industry dummies to the model of column (3). As in column (4) of Table 2, this provides us with an indication of whether the additional effect on target firms is driven by the different industrial composition of the firms in the target and non-target portfolios. We again find that our estimate is not substantially affected by this change to the specification.

While these regressions expand our identification strategy beyond the November 2013 Geneva agreement, they nonetheless still include this event. To check that our conclusions do not depend on the Geneva deal, we rerun the regression of column (3) of Table 3 for a sample period that excludes the entire sample frame of the previous subsection (the two Geneva event days and the sixty previous trading days). The results, displayed in column (1) of Table 4, are highly reassuring: our estimated coefficients are only slightly smaller in magnitude and they both remain statistically significant. Columns (2) and (3) decompose

[23]Clustering at the industry level using the wild bootstrap approach of Cameron, Gelbach and Miller (2008) yields p-values of 0.000 for our estimate of ϕ and 0.006 for our estimate of θ.

[24]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.

these coefficients into separate estimates for the periods before and after the Geneva sample period. The resulting point estimates are similar across the two periods, though our standard errors are much higher for the pre-Geneva regression (which covers a shorter period with less variation in news coverage, as is visible in Figure 2).

Finally, we reproduce the Geneva case study using the GDELT data, by running specification (2) with the sample frame we used in the previous subsection. In other words, we now include only the sample days we omitted in column (1) of Table 4.[25] The estimated effects for this period are particularly large: a one standard deviation rise in news coverage is associated with a return exceeding the mean by 0.143 percentage points for non-target firms, plus an additional 0.159 points above the mean for firms in the target portfolio.

5.2.2 Robustness

In Table 5, we assess the robustness of our estimates to variations in our definitions of the two key variables $Coverage_t$ and $Target_i$. Again, all of these results should be compared to column (3) of Table 3. In column (1) of Table 5, we use a news coverage measure that encompasses all events relating to Iran and at least one P5+1 member, instead of focusing only on events in the two-digit categories 'Express intent to cooperate', 'Consult' and 'Engage in diplomatic cooperation'. In column (2), we return to our original measure, but drop events relating to P5+1 members other than the United States, which is often perceived as the main P5+1 actor in talks with Iran. Neither change has a substantial effect on our results.

Next, rather than using a continuous measure to proxy for diplomatic progress, we return to an event study model in which 'event days' are assigned a value of one and all other days receive a value of zero. Specifically, we define only the top 10% of observations as events, which corresponds approximately to a dummy for days with a number of relevant news articles at least one standard deviation above the mean. The estimates reported in column (3) show that returns of non-target firms are 0.067 percentage points higher on these event days, and those of target firms are an additional 0.158 percentage points larger.

Columns (4) and (5) of Table 5 address the potential concern that our results for the difference-in-difference coefficient θ are sensitive to our definition of the target portfolio. We first divide our target portfolio into two subgroups: the firms identified directly from US Treasury press releases (of which six are in our sample), and the 44 additional firms identified using the Tehran Stock Exchange shareholder data (as discussed in Section 4.2). Henceforth we will call these 'Treasury firms' and 'TSE firms'. In column (4), we find that the estimated difference-in-difference coefficients are positive and statistically significant for both sets of firms, implying that our results are not unique to either of the two subsamples.

The above regression treats all TSE firms equally. However, although at least one Treasury firm is a shareholder in each of these companies, the proportion of shares held by Treasury firms may be small in some cases and large in others. We thus check the extent to which the 'low-share' TSE firms are driving our results. To do so, we define a new variable

[25] Because we omit February to August 2013 from the regressions in this section, we actually only drop 58 rather than 62 trading days in columns (1) to (3) and include only 58 trading days in column (4).

for each of the TSE firms: the proportion of shares held by Treasury firms as of the beginning of the sample period.[26] We then use this information to separate the 44 TSE firms into above-median ('high-share') and below-median ('low-share') subgroups, and augment the regression of Table 5 column (4) by estimating separate coefficients for each subgroup. Column (5) shows that it is not the low-share firms that are primarily driving our results; instead, the coefficients for both of these sets of firms are statistically significant. Moreover, the relevant point estimate is larger for high-share firms, though the difference between the two estimated coefficients is not statistically significant.

6 Alternative interpretations: conflict risk

The analysis above rests on the assumption that the content of our news coverage measure consists of 'clean' information that is predominantly about sanctions. That is, the response of market participants to news about negotiations represents their assessment of the financial implications of the changing probability of sanctions relief. A challenge to this assumption could emerge if other types of financially relevant information are conflated with sanctions-related news.

For example, the progress of negotiations has natural implications for the foreign policy environment faced by Iran, specifically the probability of military conflict between Iran and other countries. As negotiations over sanctions progressed and the probability of a successful deal increased, the likelihood of direct conflict also dropped. News about sanctions negotiations could therefore have led markets to respond by at least partly incorporating the financial implications of this easing of the probability of conflict. In econometric terms, news about sanctions could be proxying for changes in this conflict probability, biasing our estimates of the effects of sanctions relief.

We perform two exercises to assess the potential scope of this bias. We note that our estimates will be biased if two conditions are met: first, that the share prices of target and non-target firms respond to changes in conflict risk, and second, that information about conflict risk is included in our empirical measures of diplomatic progress towards sanctions removal. We therefore use additional sources of data to gather evidence on whether each of these conditions holds in practice, focusing alternatively on Iranian and non-Iranian firms.

6.1 TSE firm returns and conflict risk

Our first exercise looks at the 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 US and/or Israel could "execute an overt airstrike against Iran by December 31, 2012", as discussed in Section 4.4. We then estimate the correlation of the stock returns of firms in each portfolio with daily changes in contract prices in Table 6. To do this, we estimate the same set of specifications as in Table

[26]The denominator of this variable is the total number of shares held by all shareholders identified in the TSE shareholder data for that date.

3, now with contract price changes as the dependent variable, using data from the first day of our sample period to the end of 2012.[27]

The results indicate that increases in conflict probability, as proxied by rises in contract prices, negatively affect stock returns for firms in our sample. In the specification with firm-quarter-day-of-week fixed effects in column (3), we find that a one-dollar increase in the contract price is associated with a statistically significant fall in returns of 0.21 percentage points. However, there is no differential sensitivity between the two portfolios, since the estimated coefficient on the interaction of our target group dummy with contract prices is small and insignificant in all specifications.

We thus conclude that while the probability of military conflict is economically relevant for both non-target and target firms, the relative incidence of these effects across the two is different from the incidence that prevails for sanctions-related news. However, because conflict risk does appear to be relevant for firms across the two portfolios, it remains important that we also assess whether our measures of sanctions-related news are likely to include information about changes in this risk.

6.2 Information content of sanctions-related news measures

In our second exercise, our objective is to study the possibility that our GDELT news coverage measure and the Geneva event dummy could have conveyed confounding information about the probability of direct military conflict. In the absence of direct odds data (such as that provided by Intrade), the ideal set of firms for this exercise would be a group of firms that are 'war-sensitive' but not exposed to the effects of sanctions relief. The reaction of these war-sensitive firms to sanctions-related news coverage and the Geneva agreement would then let us gauge whether confounding information might be present over our sample period.

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 but, as non-Iranian firms, are not directly exposed to sanctions. As discussed in Section 4.4, we use a group of 66 firms from the 2012 SIPRI Arms Industry Database, which identifies the world's largest arms-producing and military services companies. Stock price information for these firms allows us to check whether they were affected by the progress of sanctions negotiations between Iran and the P5+1.

In column (1) of Table 7, we reproduce our baseline Geneva exercise using these 66 firms in the arms industry. Specifically, we define the Geneva event dummy as the day November 25, 2013 (the Monday after the deal was agreed) and add the sixty trading days prior to this to the sample.[28] We then regress the daily returns of arms industry firms on the Geneva

[27] 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 March 2013, after the US Commodity Futures Trading Commission (CFTC) applied substantial legal pressure over concerns about unregulated trading.

[28] When studying the returns of firms listed on the TSE, we defined November 23 and 24, 2013 as the two-day Geneva event. However, unlike the TSE, there was no trading on these two days (Saturday and

dummy and firm fixed effects, as in the specification in column (1) of Table 2. The results in column (1) show that the estimated coefficient on the Geneva dummy is statistically insignificant.

Similarly, column (3) displays the results of a regression based on our preferred specification in column (3) of Table 3. Using data from the full sample period, we regress daily returns on our standardized GDELT coverage measure and firm-quarter-day-of-week fixed effects. Again, the estimated coefficient of interest is not statistically significant. Along with the results in column (1), this suggests that there is not a strong relationship between our measures of progress in sanctions-related diplomatic negotiations and changes in the risk of Iran entering a war.

It is possible that we have failed to observe a statistically significant relationship in these two regressions because too few of the firms in the sample would financially benefit from increased hostilities between Iran and the West. However, this possibility is not supported by a regression using the Intrade data discussed above. A specification with daily return as the dependent variable and the change in the price of the contract "US and/or Israel execute an overt airstrike against Iran by December 31, 2012" as the regressor, along with firm fixed effects, yields an estimated coefficient of 0.301 and a standard error of 0.034. In other words, the relationship between the stock prices of these 66 firms and the risk of this particular conflict is, on average, actually very strong.

Nonetheless, we refine our exercise using a firm-specific measure of sensitivity to the odds of conflict between Iran and the US or Israel. Specifically, 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 Intrade contract prices. We then rank the firms according to the size of the estimated coefficient on contract price change (from most positive to most negative) and add the interaction between this variable and our sanctions news measure to the regressions of columns (1) and (3). If news about sanctions was interpreted by markets as a reduction in the odds of a war between Iran and the US or Israel, we should see that the returns of firms who are more sensitive to such a conflict fall further when this news emerges.

Column (2) displays the results of a regression adding the interaction of the Geneva dummy and this sensitivity ranking to the specification of column (1), while column (4) adds the interaction between the GDELT measure and war sensitivity to the regression of column (3). In both cases, the estimated coefficient on the interaction term is statistically insignificant. There is thus no evidence that firms that are more sensitive to the odds of conflict between Iran and other countries respond differently to news shocks about the diplomatic negotiation process.

In conclusion, the evidence from these two exercises does not support a confounding role for the risk of direct conflict in our estimates. The first exercise suggests that target portfolio firms are not differentially sensitive to conflict risk, while the second indicates that, in any case, updates to expectations of conflict risk were not linked to progress in sanctions negotiations.

Sunday) on the stock markets considered in this exercise.

7 Conclusion

In this study, we have sought to test a central plank of modern international sanctions policy: the efficacy of the targeting of particular actors within a receiver country. We consider the case of Iran, examining responses of its stock market to information indicating progress in diplomatic negotiations towards an agreement on sanctions removal. We find evidence that firms owned by key groups within Iran's political system have been differentially sensitive to news about potential sanctions relief. Importantly, this sensitivity suggests that political elites within Iran faced a tangible economic incentive for the removal of sanctions.

Our results therefore provide evidence that sanctions were at least partly 'on target' in the sense that elite policymakers appear to have benefited from their removal. Note that in the case where this elite 'income targeting' goal is not met, then sanctions can still work as an incentive for policy reform. However, in that case they need to impose costs more heavily on non-elite groups, a 'worst-case scenario' for sanctions policy as it is currently framed by policymakers. The fact that we identify that an economic incentive for sanctions removal existed for politically influential groups suggests that this loomed as a potential input into Iran's political decision-making process regarding its nuclear programme.

References

[1] Acemoglu, Daron, Tarek A. Hassan and Ahmed Tahoun (forthcoming). "The Power of the Street: Evidence from Egypt's Arab Spring", *Review of Financial Studies*.

[2] Acemoglu, Daron, Simon Johnson, Amir Kermani, James Kwak and Todd Mitton (2016). "The Value of Connections in Turbulent Times: Evidence from the United States", *Journal of Financial Economics* 121 (2), 368-391.

[3] Ahn, Daniel P. and Rodney D. Ludema (forthcoming). "Measuring Smartness: The Economic Impact of Targeted Sanctions against Russia", in Tibor Besedes and Volker Nitsch (eds.), *Disrupted Economic Relationships: Disasters, Sanctions, Dissolutions*, MIT Press.

[4] Alfoneh, Ali (2010). "The Revolutionary Guards' Looting of Iran's Economy", *Middle Eastern Outlook* No. 3, Washington, D.C.: American Enterprise Institute.

[5] Allen, Susan Hannah (2008). "The Domestic Political Costs of Economic Sanctions", *Journal of Conflict Resolution* 52 (6), 916-944.

[6] BBC (2012). "Iran nuclear talks in Istanbul 'constructive"', http://www.bbc.com/news/world-middle-east-17716241.

[7] Berman, Eli, Jacob N. Shapiro and Joseph H. Felter (2011). "Can Hearts and Minds be Bought? The Economics of Counterinsurgency in Iraq", *Journal of Political Economy* 119 (4), 766-819.

[8] Braggion, Fabio and Lyndon Moore (2013). "The Economic Benefits of Political Connections in Late Victorian Britain", *Journal of Economic History* 73 (1), 142-176.

[9] Cameron, A. Colin, Jonah B. Gelbach and Douglas L. Miller (2008). "Bootstrap-based Improvements for Inference with Clustered Errors", *Review of Economics and Statistics* 90 (3), 414-427.

[10] Čihák, Martin, Aslı Demirgüç-Kunt, Erik Feyen and Ross Levine (2012). "Benchmarking Financial Systems Around the World", World Bank Policy Research Paper 6175.

[11] Coulomb, Renaud and Marc Sangnier (2014). "The Impact of Political Majorities on Firm Value: Do Electoral Promises or Friendship Connections Matter?", *Journal of Public Economics* 115, 158-170.

[12] DellaVigna, Stefano and Eliana La Ferrara (2010). "Detecting Illegal Arms Trade", *American Economic Journal: Economic Policy* 2 (4), 26-57.

[13] Dube, Arindrajit, Ethan Kaplan and Suresh Naidu (2011). "Coups, Corporations and Classified Information", *Quarterly Journal of Economics* 126 (3), 1375-1409.

[14] Durante, Ruben and Ekaterina Zhuravskaya (forthcoming). "Attack When the World is Not Watching: U.S. Media and the Israeli-Palestinian Conflict", *Journal of Political Economy*.

[15] Eaton, Jonathan and Maxim Engers (1992). "Sanctions", *Journal of Political Economy* 100 (5), 899-928.

[16] European Commission (2008). "Restrictive measures", http://eeas.europa.eu/cfsp/sanctions/docs/index_en.pdf.

[17] Faccio, Mara (2006). "Politically Connected Firms", *American Economic Review* 96 (1), 369-386.

[18] Fisman, Raymond (2001). "Estimating the Value of Political Connections", *American Economic Review* 91 (4), 1095-1102.

[19] Guidolin, Massimo and Eliana La Ferrara (2007). "Diamonds Are Forever, Wars Are Not: Is Conflict Bad for Private Firms?", *American Economic Review* 97 (5), 1978-1993.

[20] Haidar, Jamal Ibrahim (2017). "Sanctions and Export Deflection: Evidence from Iran", *Economic Policy* 32 (90), 319-355.

[21] Hsieh, Chang-Tai and Enrico Moretti (2006). "Did Iraq Cheat the United Nations? Underpricing, Bribes and the Oil for Food Program", *Quarterly Journal of Economics* 121 (4), 1211-1248.

[22] Hufbauer, Gary, Jeffrey Schott and Kimberley Elliott (1990). *Economic Sanctions Revisited: Supplemental Case Histories*, Washington, D.C.: Institute for International Economics.

[23] Jaeger, David A. and M. Daniele Paserman (2006). "Israel, the Palestinian Factions, and the Cycle of Violence", *American Economic Review: Papers and Proceedings* 96(2), 45-49.

[24] Jaeger, David A. and M. Daniele Paserman (2008). "The Cycle of Violence? An Empirical Analysis of Fatalities in the Palestinian-Israeli Conflict", *American Economic Review* 98(3), 1591-1604.

[25] Jaeger, David A. and M. Daniele Paserman (2009). "The Shape of Things to Come? Assessing the Effectiveness of Suicide Attacks and Targeted Killings", *Quarterly Journal of Political Science* 4(4), 315-342.

[26] Jaeger, David A., Esteban F. Klor, Sami H. Miaari and M. Daniele Paserman (2012). "The Struggle for Palestinian Hearts and Minds: Violence and Public Opinion in the Second Intifada", *Journal of Public Economics* 96 (3-4), 354-368.

[27] Jahan-Parvar, Mohammad R. and Hassan Mohammadi (2013). "Risk and Return in the Tehran Stock Exchange", *Quarterly Review of Economics and Finance* 53 (3), 238-256.

[28] Kaempfer, William H. and Anton D. Lowenberg (1988). "The Theory of International Sanctions: A Public Choice Approach", *American Economic Review* 78 (4), 786-793.

[29] Katzman, Kenneth and Paul K. Kerr (2017). "Iran Nuclear Agreement", Congressional Research Service Report R43333.

[30] Kerr, Paul K. (2017). "Iran's Nuclear Program: Status", Congressional Research Service Report RL34544.

[31] MacKinley, A. Craig (1997). "Event Studies in Economics and Finance", *Journal of Economic Literature* 35 (1), 13-39.

[32] Marinov, Nikolay (2005). "Do Economic Sanctions Destabilize Country Leaders?", *American Journal of Political Science* 49 (3), 564-576.

[33] Nasim Online (2013). "What is Happening in Third Round of Nuclear Talks in Geneva", http://en.nasimonline.ir/archives/11665.

[34] Neuenkirch, Matthias and Florian Neumeier (2016). "The Impact of US Sanctions on Poverty", *Journal of Development Economics* 121 (1), 110-119.

[35] Reuters (2012) "Iran: Military Decides if IAEA can visit Parchin", http://www.reuters.com/article/us-nuclear-iran-parchin-idUSTRE81S0S820120229.

[36] Reuters (2013). "Assets of the Ayatollah: The Economic Empire Behind Iran's Supreme Leader", http://www.reuters.com/investigates/iran/.

[37] Reuters (2015). "Conglomerate Controlled by Iran's Supreme Leader a Winner from Nuclear Deal", http://www.reuters.com/article/us-iran-nuclear-setad-insight-idUSKCN0RF0E920151021.

[38] Schrodt, Philip A. (2012). "CAMEO Conflict and Mediation Event Observations Event and Actor Codebook", http://data.gdeltproject.org/documentation/CAMEO.Manual.1.1b3.pdf.

[39] Tehran Stock Exchange (2011). "Iran's Capital Market: How to Invest in Tehran Stock Exchange (TSE)", mimeograph.

[40] United States Treasury (2013). "Treasury Targets Assets of Iranian Leadership", press release, June 4 2013, https://www.treasury.gov/press-center/press-releases/Pages/jl1968.aspx.

[41] Wall Street Journal (2006). "A Feared Force Roils Business in Iran: As Hard-Liners Rise, Shadowy Revolutionary Guard Muscles in on Airport and Nabs Energy Deals", http://www.wsj.com/articles/SB116079428301492832.

Figures and Tables

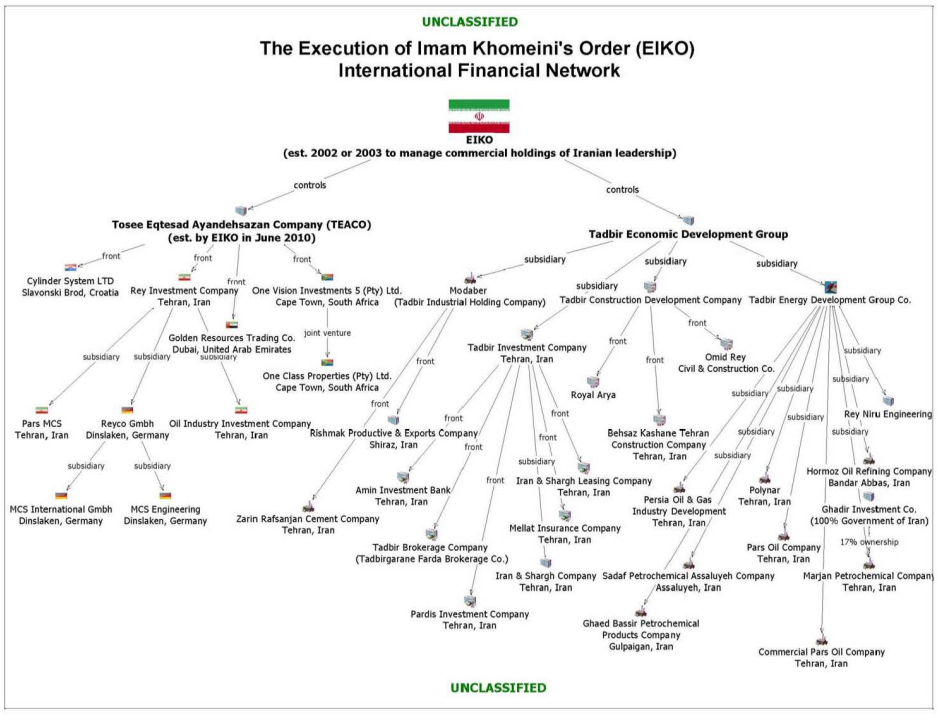

Figure 1: Setad ownership diagram (from US Treasury)

This figure reproduces an ownership diagram for Setad (also known as EIKO) published by the US Treasury in a 2013 press release.

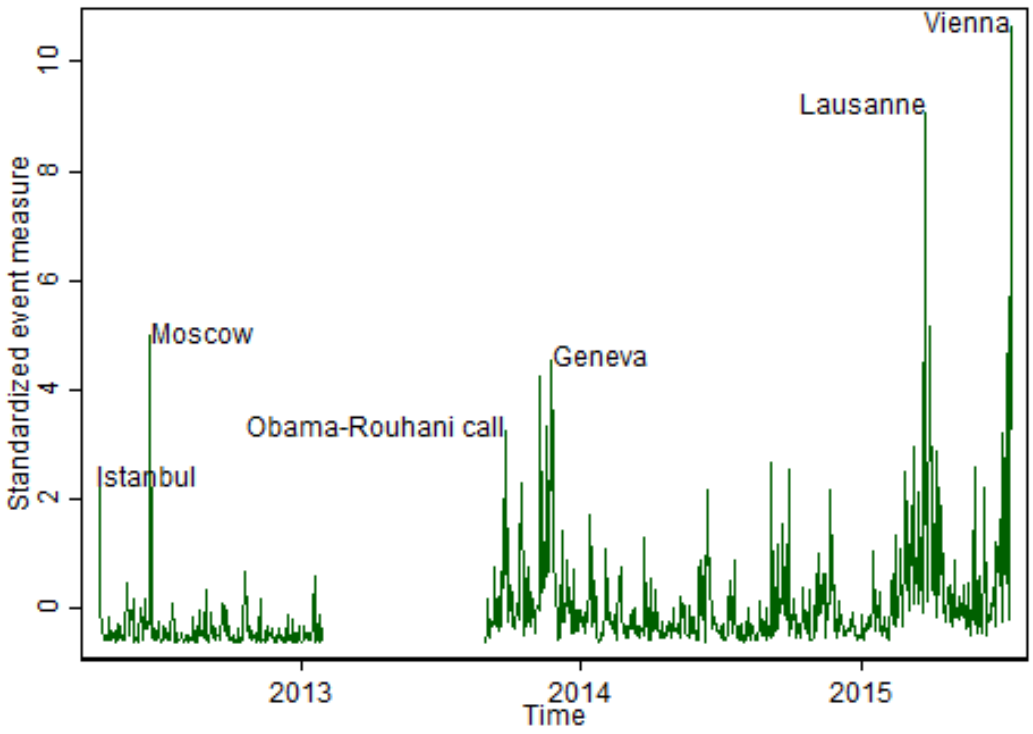

Figure 2: Value of daily event measure, April 14, 2012 to July 15, 2015

This figure shows the evolution, over the sample period April 14, 2012 to July 15, 2015 (excluding February to August 2013), of a standardized count of the number of articles on relevant events identified in the GDELT dataset, as discussed in Section 4.3.

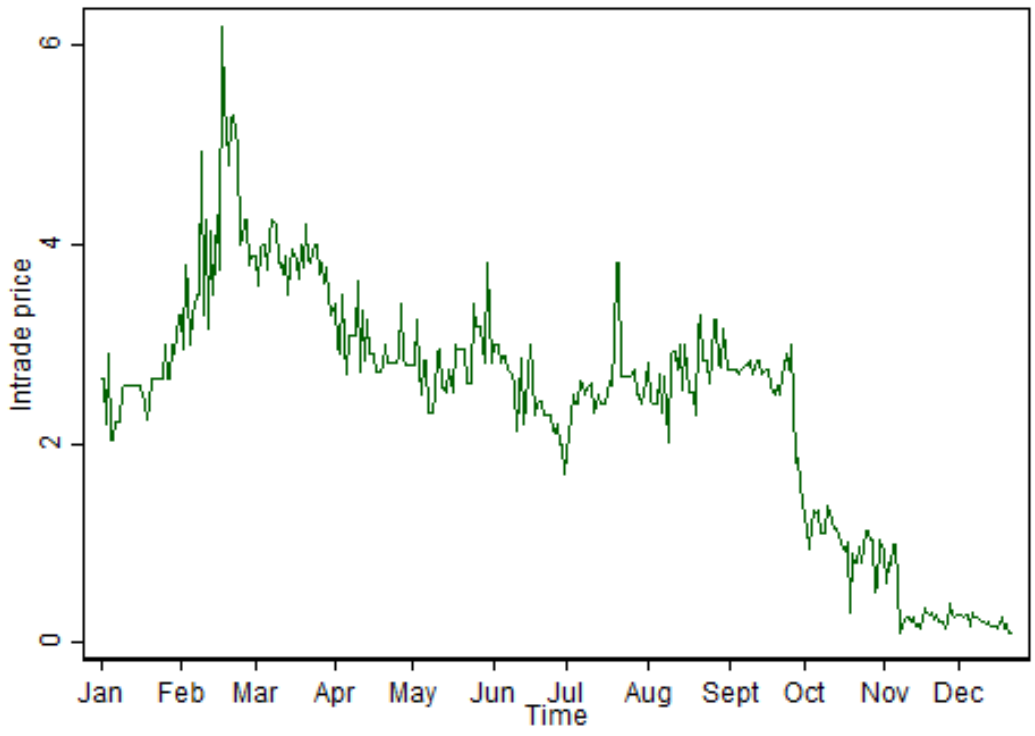

Figure 3: Daily price of Intrade contract in 2012

This figure shows the evolution of Intrade contract prices (in current US dollars) for the betting contract 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. The period plotted is January 1, 2012 to December 22, 2012 (the day before US trading was suspended).

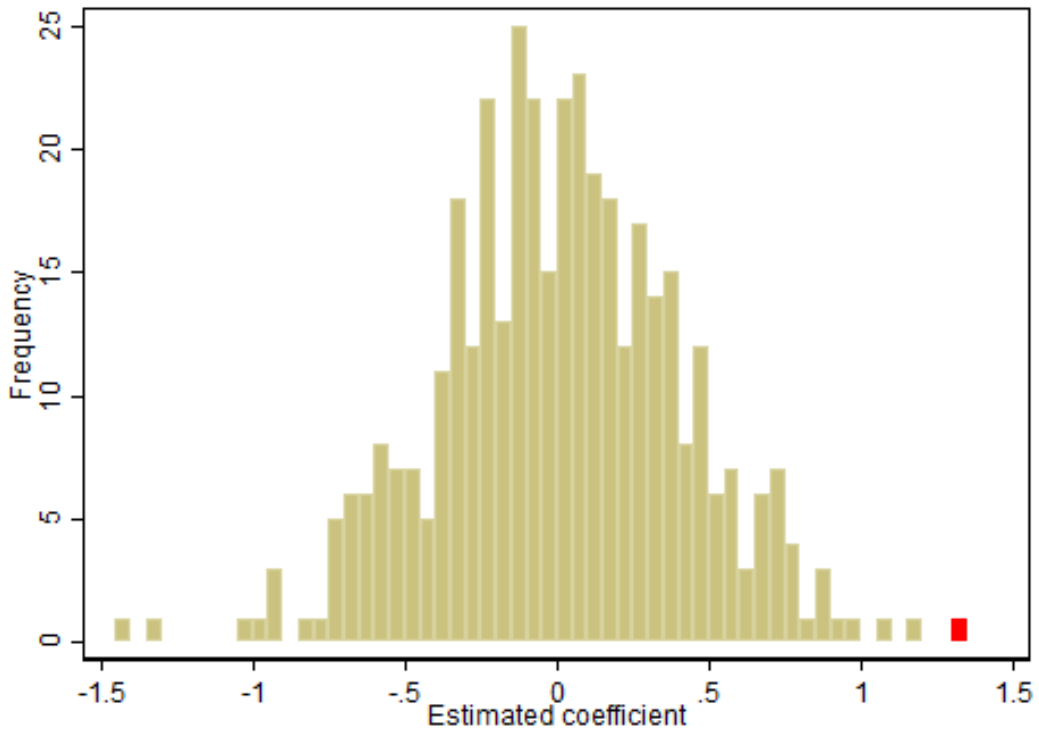

Figure 4: Distribution of estimates across two-day periods (Geneva deal highlighted)

In this figure, we use data for the period from April 14-15, 2012 until July 14-15, 2015 and estimate a series of models with different 2-day dummies interacted with our target group definition. This corresponds to 388 regressions, and the histogram reports the distribution of estimated coefficients with the two-day Geneva event (November 23-24, 2013) marked in red.

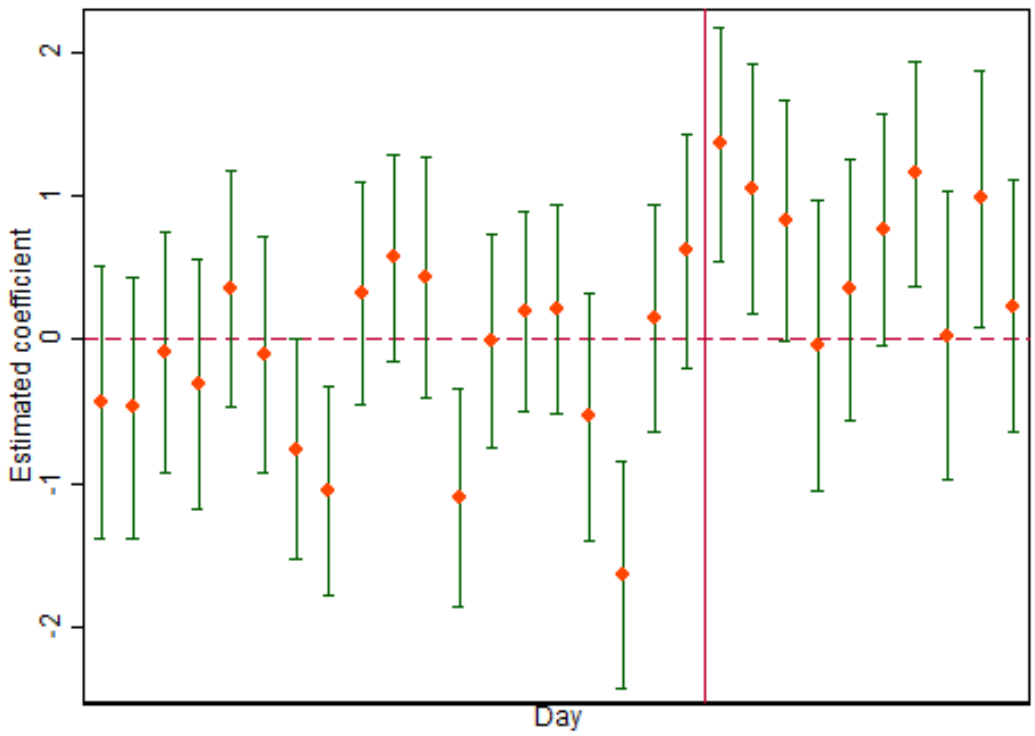

Figure 5: Daily estimated target-day interaction coefficients

In this figure, we present estimates of daily target * day interactions for the last four weeks of the 60-day estimation window in Table 2, column (1), as well as the two trading weeks starting November 23, 2013 (which are to the right of the vertical line). Estimated coefficients and 95% confidence intervals are plotted. Firm fixed effects are included and standard errors are clustered by firm.

Table 1: Summary statistics – sample firms

	Target firms	Non-target firms
Number of firms	50	87
Daily return	0.134	0.135
	(2.01)	(2.06)
Market capitalization	12,357.47	7,692.31
	(31,273.68)	(22,499.35)
Share of TSE market capitalization	21.53%	23.32%
Industry of firm:		
Mining	1	4
Food products	2	6
Wood/paper/textiles	1	3
Coke/refined petroleum	2	1
Chemicals	1	8
Pharmaceuticals	3	9
Rubber/plastic/mineral products	14	16
Basic metals	5	9
Metal products	1	1
Electronics/electrical equipment	2	7
Machinery	2	2
Motor vehicles	6	7
Transportation/telecom	1	1
Finance	8	7
Construction/real estate	1	6
Orbis data:		
Number of firms	32	56
Turnover	3,566.48	4,042.20
	(8,899.41)	(23,957.31)
Assets	5,674.74	3,310.79
	(16,403.98)	(9,581.07)
Employees	1,679.71	778.92
	(3,451.24)	(1,281.56)

This table displays summary statistics for target and non-target firms. See Section 4.1 for a discussion of the industrial classification. Daily stock return is in percent; its mean and standard deviation (in parentheses) are by firm-day, omitting the top and bottom 1% of observed returns. Market capitalization, turnover and assets are denominated in billion rial. Market capitalization from TSE data is as of March 10, 2014. Total TSE market capitalization is calculated for firms listed on the TSE as of April 14, 2012, omitting two firms for which data is unavailable. Data from Orbis on turnover, assets and employees is from 2012, and is 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 is missing for an additional nine firms (four target firms and five non-target firms).

Table 2: Results – Geneva deal

	(1) Firm FEs	(2) Day-by-day estimates	(3) Firm-day-of-week FEs	(4) Industry interactions
Geneva	0.610**		0.699**	
	(0.279)		(0.289)	
Day 1		0.521*		
		(0.307)		
Day 2		0.700**		
		(0.344)		
Target * Geneva	1.294***		1.244***	1.274***
	(0.366)		(0.386)	(0.400)
Target * Day 1		1.442***		
		(0.419)		
Target * Day 2		1.147**		
		(0.447)		
Firm FEs	Yes	Yes	No	Yes
Firm-day-of-week FEs	No	No	Yes	No
Observations	6,528	6,528	6,528	6,528
Number of firms	128	128	128	128

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 percent. '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. Columns (1), (2) and (4) include firm fixed effects, and column (3) includes firm-day-of-week fixed effects. Column (4) includes industry * Geneva interactions for 15 industries. Standard errors, clustered by firm, are in parentheses. Small p-values are represented by *** (below 1%), ** (below 5%) or * (below 10%).

Table 3: Results – news coverage

	(1) Firm FEs	(2) Firm-quarter FEs	(3) Firm-quarter-day-of-week FEs	(4) Industry interactions
Coverage	0.068***	0.066***	0.063***	
	(0.011)	(0.011)	(0.013)	
Target * coverage	0.066***	0.056**	0.069***	0.075***
	(0.016)	(0.016)	(0.017)	(0.018)
Firm FEs	Yes	No	No	No
Firm-qtr FEs	No	Yes	No	No
Firm-qtr-day FEs	No	No	Yes	Yes
Observations	61,568	61,568	61,568	61,568
Number of firms	137	137	137	137

This table displays estimated effects 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 percent. The variable 'coverage' is a standardized count of the number of articles on a relevant event identified in the GDELT dataset, as discussed in Section 4.3. Column (1) includes firm fixed effects, column (2) includes firm-quarter fixed effects and columns (3) and (4) include firm-quarter-day-of-week fixed effects. Column (4) includes industry * coverage interactions for 15 industries. The sample period is from April 14, 2012 to July 15, 2015, excluding February to August 2013. Standard errors, clustered by firm, are in parentheses. Small p-values are represented by *** (below 1%), ** (below 5%) or * (below 10%).

Table 4: Results – varying sample period

	(1) No Geneva	(2) Pre-Geneva	(3) Post-Geneva	(4) Geneva
Coverage	0.061***	0.043	0.063***	0.143***
	(0.012)	(0.039)	(0.013)	(0.038)
Target * coverage	0.046**	0.028	0.049**	0.159**
	(0.018)	(0.053)	(0.020)	(0.051)
Firm-qtr-day FEs	Yes	Yes	Yes	Yes
Observations	55,441	16,122	39,319	6,127
Number of firms	137	137	130	128

This table displays estimated effects 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 percent. The variable 'coverage' is a standardized count of the number of articles on a relevant event identified in the GDELT dataset, as discussed in Section 4.3. The sample period in column (1) is from April 14, 2012 to July 15, 2015 excluding February 1, 2013 to November 24, 2013, in column (2) it is from April 14, 2012 to January 31, 2013, in column (3) it is from November 25, 2013 to July 15, 2015 and in column (4) it is from September 1, 2013 to November 24, 2013. All columns include firm-quarter-day-of-week fixed effects. Standard errors, clustered by firm, are in parentheses. Small p-values are represented by *** (below 1%), ** (below 5%) or * (below 10%).

Table 5: Robustness – varying coverage and target definitions

	(1) All articles	(2) US-Iran only	(3) Top 10% dummy	(4) 'Target' by data source	(5) High-share and low-share firms
Coverage	0.049***	0.055***	0.067**	0.063***	0.066***
	(0.013)	(0.013)	(0.033)	(0.013)	(0.012)
Target * coverage	0.069***	0.083***	0.158***		
	(0.017)	(0.018)	(0.050)		
US Treasury * coverage				0.073*	0.070*
				(0.042)	(0.041)
TSE asset * coverage				0.068***	
				(0.018)	
High share * coverage					0.090***
					(0.029)
Low share * coverage					0.064**
					(0.028)
Firm-qtr-day FEs	Yes	Yes	Yes	Yes	Yes
Observations	61,568	61,568	61,568	61,568	61,568
Number of firms	137	137	137	137	137

This table displays estimated effects 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 percent. In all columns, the variable 'coverage' is based on the number of articles on a event identified in the GDELT dataset. In column (1), this includes all events, while in all other columns, it is restricted to relevant events as discussed in Section 4.3. In column (2), this includes only events involving the US, while in all other columns, it includes events involving any P5+1 country. The coverage variable is a dummy for the top 10% of observed values in column (3) and a standardized measure in all other columns. Column (4) classifies target firms into two groups: firms identified by the US Treasury as assets of IRGC and Setad ('US Treasury') and firms identified via the TSE shareholder data as assets of US Treasury-identified firms. Column (5) additionally classifies firms identified via the TSE shareholder data into two groups: those with above-median ownership by US Treasury-identified firms ('high share') and those with below-median ownership by US Treasury-identified firms ('low share'). All columns include firm-quarter-day-of-week fixed effects. The sample period is from April 14, 2012 to July 15, 2015, excluding February to August 2013. Standard errors, clustered by firm, are in parentheses. Small p-values are represented by *** (below 1%), ** (below 5%) or * (below 10%).

Table 6: Results – Intrade contract price

	(1) Firm FEs	(2) Firm-quarter FEs	(3) Firm-quarter-day-of-week FEs	(4) Industry interactions
Price change	-0.160***	-0.183***	-0.209***	
	(0.060)	(0.061)	(0.070)	
Target * price change	0.036	0.027	0.015	-0.075
	(0.092)	(0.093)	(0.103)	(0.103)
Firm FEs	Yes	No	No	No
Firm-qtr FEs	No	Yes	No	No
Firm-qtr-day FEs	No	No	Yes	Yes
Observations	12,332	12,332	12,332	12,332
Number of firms	136	136	136	136

This table displays correlations between returns of target and non-target firms and the change in the price of the Intrade contract "US and/or Israel to execute an overt airstrike against Iran by December 31, 2012". The dependent variable is daily stock return in percent. 'Price change' represents the change in the price of the Intrade contract, in current US dollars, during the same day. Column (1) includes firm fixed effects, column (2) includes firm-quarter fixed effects and columns (3) and (4) include firm-quarter-day-of-week fixed effects. Column (4) includes industry * coverage interactions for 15 industries. The sample period is from April 14, 2012 to December 22, 2012. Standard errors, clustered by firm, are in parentheses. Small p-values are represented by *** (below 1%), ** (below 5%) or * (below 10%).

Table 7: Results – arms industry firms

	(1)	(2)	(3)	(4)
	Geneva deal		News coverage	
Geneva	-0.061	-0.157		
	(0.140)	(0.257)		
Coverage			0.008	0.011
			(0.006)	(0.014)
Sensitivity rank * Geneva		0.003		
		(0.007)		
Sensitivity rank * coverage				-0.00008
				(0.0003)
Firm FEs	Yes	Yes	No	No
Firm-qtr-day FEs	No	No	Yes	Yes
Observations	3,977	3,977	45,040	45,040
Number of firms	66	66	66	66

This table displays estimated effects of the Geneva deal, and estimated effects from specifications that include a daily measure of news coverage related to diplomatic progress between Iran and the P5+1 countries, on returns of firms in the arms industry. The dependent variable is daily stock return in percent. 'Geneva' is defined as the day Monday November 25. The variable 'coverage' is a standardized count of the number of articles on a relevant event identified in the GDELT dataset, as discussed in Section 4.3. The variable 'sensitivity rank' is the rank of each firm based on the responsiveness of its stock return to the change in the price of the Intrade contract "US and/or Israel to execute an overt airstrike against Iran by December 31, 2012" in the first quarter of 2012, as discussed in Section 6.2. Columns (1) and (2) include firm fixed effects and columns (3) and (4) include firm-quarter-day-of-week fixed effects. The sample period in columns (1) and (2) is Geneva and the previous sixty trading days, and in columns (3) and (4) it is April 16, 2012 to July 15, 2015. Standard errors, clustered by firm, are in parentheses. Small p-values are represented by *** (below 1%), ** (below 5%) or * (below 10%).

A1 Appendix Tables

Table A1: Two-digit CAMEO event types

Code	Event type
01	Make public statement
02	Appeal
03	Express intent to cooperate
04	Consult
05	Engage in diplomatic cooperation
06	Engage in material cooperation
07	Provide aid
08	Yield
09	Investigate
10	Demand
11	Disapprove
12	Reject
13	Threaten
14	Protest
15	Exhibit military posture
16	Reduce relations
17	Coerce
18	Assault
19	Fight
20	Engage in unconventional mass violence

This table displays the event types at the two-digit level of the Conflict and Mediation Event Observations (CAMEO) classification system (Schrodt 2012).

Table A2: Top fifteen event measure observations

Date	Value	Sanctions negotiations event
Jul 14 2015	10.66	Final agreement reached in Vienna
Mar 25 2015*	9.07	Negotiation round ends in Lausanne
Apr 4 2015*	5.16	Framework agreement reached in Lausanne
Jun 19 2012*	4.97	Negotiation round ends in Moscow
Jul 11 2015*	4.67	Negotiation round in progress in Vienna
Nov 24 2013	4.55	Interim agreement reached in Geneva
Nov 9 2013*	4.22	Negotiation round ends in Geneva
Nov 23 2013*	3.55	Negotiation round in progress in Geneva
Nov 18 2013	3.33	Negotiation round about to begin in Geneva
Jul 15 2015	3.26	Final agreement reached on previous day in Vienna
Sep 28 2013*	3.23	Presidents of Iran and US speak by phone
Jul 4 2015*	3.21	Negotiation round in progress in Vienna
Mar 31 2015	3.17	Negotiation deadline extended in Lausanne
Mar 7 2015*	2.95	Several P5+1 members discuss negotiations
Mar 28 2015*	2.89	Negotiation round in progress in Lausanne

This table displays the top fifteen observations, within the sample period April 14, 2012 to July 15, 2015, of a standardized count of the number of articles on relevant events identified in the GDELT dataset, as discussed in Section 4.3. For each of these observations, the table shows the relevant date, the value of the variable, and an event of that day that is relevant to diplomatic negotiations for sanctions relief. Dates are marked with stars when the observation also includes events from prior days because of weekends or holidays on the TSE.