ERASMUS UNIVERSITY ROTTERDAM

Erasmus School of Economics

Master Thesis in International Economics

Analysing international transmission of shocks related to the 2022 Russian invasion of Ukraine using transaction-level data on online business-to-business sales

Abstract

Major geopolitical events usually have a profound impact on the global economy with the 2022 war in Ukraine being the most recent and striking example. Studying them using real-time data can help understand their consequences and inform policy responses. In this thesis, I analyse how online business-to-business sales evolved in the immediate aftermath of the war and how that reaction depended on a country's exposure to the conflict. I use a panel dataset of over 7,5 million transactions recorded with daily frequency and measure changes in the transactions' number and value. My results suggest that, on average, more transactions took place per day during the war than before and their average daily value was higher. Moreover, highly-exposed countries experienced smaller changes compared to the rest of the sample. These unexpected results provide a valuable starting point in understanding the effects of the war in Ukraine on economic outcomes in countries not involved in it directly.

Student: Michalina Lesińska

Student number: 576331ml

Supervisor: dr. (Sacha) SV Kapoor

Second assessor: dr. (Felix) FPL Ward

Date of the final version: 26.07.2022

Table of contents

1. Introduction

With the interconnectedness that characterises the modern economy, events happening in one part of the globe might greatly affect economic outcomes worldwide, often in unexpected ways. The 2022 Russian invasion of Ukraine is a perfect example, as it was largely unforeseen and continues to have a major impact even on countries not involved in the war or any of the economic sanctions. Critical sectors such as food and energy production experience profound disruptions putting at risk the livelihoods of hundreds of millions of people.

The consequences of the war in Ukraine are still unfolding as are events on the battlefield, resulting in a grim future outlook for the global economy. At the same time, even the immediate, short-term impact of the invasion is yet to be fully understood and the picture is complex due to the spillover effects interacting with each other. Understanding their relative importance is crucial for formulating appropriate policy responses, but there is little empirical evidence within this domain as the war is a recent event. With numerous factors potentially playing a role, economic theory cannot provide definitive answers, therefore turning our efforts to empirics and exploring the available data is necessary to guide a more rigorous analysis. Without data on transactions, such as that I have access to, precise and meaningful analysis is extremely difficult. Lastly, a better understanding of the real-time behaviour of transactions can also lend insights into future aggregate outcomes.

In this thesis, I use new and proprietary data from online business-to-business stores to measure the economic effects of the 2022 war in Ukraine in nearly real-time. With the focus on business-to-business relations, I can provide novel insights into an area of economic activity which usually does not receive much scholarly attention from economists. The broad goal is to advance understanding of the impact of geopolitical shocks on demand and the role of supply chain linkages in their propagation. Changes brought about by the war, such as higher uncertainty, rising transportation costs and unprecedented policies could potentially depress demand yet the impact on different sectors and firms might vary greatly.

I have access to a panel dataset of online sales of 536 firms from 43 countries, spanning 27 months, from the 11th of February 2020 to the 11th of April 2022, with daily frequency. I focus on the daily pattern of both the number of transactions and their value for each firm and currency of the transactions. These are used as dependent variables for the regressions run in this thesis. The data was collected by Sana Commerce, an e-commerce software provider. I supplement it with two country-level exposure measures that proxy for the extent to which economic agents in the country of the buyers are potentially impacted by the war. These are the importance of trade with Russia, Belarus and Ukraine, and the dependence on Russian gas. Moreover, I use monthly data on foreign exchange rates and the Consumer Price Indexes to be able to compare the daily values of transactions across countries and time.

The main analysis consists of three parts. First, I estimate regressions at the original, firm-level aggregation where I include a single dummy variable to indicate days after the beginning of the war, one exposure measure at the time and the interaction between the time dummy indicating the war and the exposure measure. Next, I perform an analogous analysis using data aggregated to the country level, depending on the country of the buyers. Lastly, I carry out an event study on the firm-level related to work such as Dunn, Hood, Batch, and Driessen (2020). In the event study, the regression equations include a separate dummy variable for each day from the 14th of February 2022 to the end of the sample period in an attempt to precisely trace the pattern of changes in the dependent variables.

The primary finding of this thesis is that the period during the war in Ukraine is robustly associated with increases in both the daily value and the daily number of online business-to-business transactions. Their size is economically significant, with the number of transactions increasing by between 17% to 25%, on average, and the value of transactions increasing by around 15% to 28%. Moreover, my results suggest that the transactions' number and the transactions' value react less strongly in countries that are potentially highly exposed to the war. This result holds for both the trade exposure measure and the gas dependence and, remarkably, both deliver quantitatively similar estimates.

Via the event study, I provide evidence that the first days immediately after the start of the war were not associated with dramatic changes in the patterns of online business-to-business transactions. Instead, the biggest increases in both dependent variables compared to the pre-war period occurred towards the end of the sample period. Further analysis reveals that most of the baseline results are robust to the country-level aggregation of data along the cross-sectional dimension and to both weekly and monthly aggregation along the time series dimension. Finally, a version of the baseline regression is run separately for every country in the sample which highlights substantial differences between countries but suggests no clear geographical pattern.

There are reasons to believe that my findings are not generalisable to the wider economy. Nonetheless, they offer a fresh perspective on the broad and important topic of economic shocks' propagation. My work falls within a relatively new branch of research which uses high-frequency, highly disaggregated and real-time data to address macroeconomic questions. Such data is usually private information of firms, making it exceedingly difficult to obtain. The businesses-to-business nature of the transactions I study and focusing on the very recent war in Ukraine leave this thesis without close counterparts in the literature. However, there have been a number of papers that use a similar approach to study other events, notably the Covid-19 pandemic or major natural disasters, which this thesis takes inspiration from.

2. Literature Review

There exists an extensive literature on the transmission and amplification of macroeconomic shocks and the various channels that can enable it. As has been established in empirical work, supply chain linkages constitute an important vector for shock propagation. To provide theoretical foundations to understand this relationship, Johnson (2014) develops a multicountry model and, by estimating regressions on simulated data, investigates the strength of the relationship between trade in intermediate inputs and comovement in output between countries. One of his conclusions is that business cycle comovement for goods is high when the aggregate trade elasticity is low. Trade elasticity captures how easy it is to substitute between domestic and foreign outputs and is, therefore, a key variable determining the extent to which crises such as the war in Ukraine will affect aggregate output abroad. Intuitively, if it is difficult to replace goods from Ukraine, Belarus and Russia with some goods produced elsewhere, the impact of the war on the output in other countries will be more severe. This also implies substantial differences between industries depending on how important the countries involved in the war are for the production of the industry's inputs. However, the author also emphasises that if the output of services is weakly correlated across countries, the aggregate comovement of outputs will be small.

In general, empirical research that is closely related to my work focuses on a particular exogenous event that disrupts economic activity and attempts to measure its consequences. Boehm, Flaaen, and Pandalai-Nayar (2019) investigate the transmission of shocks caused by the Tōhoku Earthquake, which hit Japan in 2011, via supply chain linkages. Looking at transactions between firms, the authors estimate that the short-run elasticity of substitution between different inputs is close to 0. This implies that for firms in other countries it is difficult to switch suppliers and avoid being impacted by the shock originating in Japan. In the light of this result, I would expect that firms that source inputs from Ukraine, Russia and

Belarus, will experience a drop in output after the war started. Moreover, the magnitude of the decrease should be proportional to how much each firm trades with firms from countries involved in the war.

Recently, the Covid-19 pandemic motivated extensive research into the propagation of demand shocks and their short-term impact on economic activity. McCann and Myers (2020), for the case of Ireland, document in rich detail how lockdown policies reduce the revenue of directly-impacted sectors (for example hotels) but also, crucially, how some non-directly affected sectors (for example suppliers of the hotels) were hit due to spillover effects. They highlight that the nature of the initial pandemic shock, which drastically affected the behaviour of final consumers, meant that transmission of the demand shock from downstream firms to upstream suppliers was more important than in the opposite direction. The upstream suppliers tended to be affected indirectly. One can reason that likely the opposite was true when the war in Ukraine started and that the upstream suppliers, who are the main focus of this thesis, were affected more strongly than business-to-consumer firms. This holds as long as intermediate inputs are more important than final goods in exports of countries involved in the war.

Lastly, Dunn et al (2020) is a paper that is particularly closely related to my work. Using near real-time daily data on card transactions, the authors assess the impact of the Covid-19 pandemic on consumer spending across a number of sectors. They are able to identify the effects of almost contemporary events with high precision and measure sector-specific as well as aggregate changes. For example, they provide evidence that sales in aggregate retail and food services increased in the first days after the pandemic was declared on the 11th of March 2022. This was followed by a sharp decrease starting in the middle of March. Analysis of this type - the event study - is partially replicated in this thesis.

Considering the above discussion, I can expect that when the 2022 war in Ukraine erupted, the reduction in output in countries involved in the war, as well as economic sanctions imposed by third countries, led to a reduction in the number and the value of business-to-business transactions around the world. This is because the war sharply decreased the amount of output produced in affected economies, especially in Ukraine, impacting firms that relied on it for inputs. Moreover, for a large portion of outputs that were produced in Russia and Belarus, the sanctions effectively blocked their movement via trade.

Further, I hypothesise that there is a positive relationship between the economic closeness of a country to Russia, Belarus and Ukraine and the degree of the abovementioned decrease in the transactions number and value. Lastly, I anticipate that the first day of the war is not the only relevant cut-off point to detect a change in buyers' behaviour. A more important but harder to pinpoint one could be the moment when it became apparent that the severity of Western sanctions exceeded expectations and that the war could continue for longer than thought.

In addition to the trade in inputs channel, there might have been other, arguably more powerful, factors at play. An important one is an increase in trade costs caused by disruption of shipping routes and growing prices of fuel. With higher trade costs, exporting becomes unprofitable for the least productive exporters and the volume of exports goes down (Melitz, 2003). This could be another channel through which the war in Ukraine disrupted supply chains and potentially resulted in a lower number and value of transactions due to the unavailability of products that buyers would be willing to purchase.

In addition to the immediate, supply-side effects of the war which were related to trade disruption, firms experienced a broader change in their business environment driven by a shift in expectations and increased uncertainty. This is illustrated by the Consumer Confidence Indicator which for the European Union dropped from -11.7 in February 2022 to -20.8 in

March (European Statistical Office, 2022a). For the same period and countries, the Industrial Confidence Indicator, particularly relevant for the sample of firms studied in this thesis, declined from 12.4 to 7.9. Crucially, the war has added to expectations of inflation growth.

Other changes that impacted firms after the beginning of the war include more serious cybersecurity threats or reputational risks if they continued operating in Russia or trading with Russian firms. Overall, this paints a highly complex picture and informs the decision that identifying the channel of the studied relationship is beyond the scope of this thesis.

3. Data

3.1. Main data source and resulting limitations

I have access to a unique, proprietary dataset of transactions from Sana Commerce, a software company which provides an e-commerce software product designed for business-to-business use. Sana Commerce is a software vendor, not a platform provider like for example Amazon. Every customer of Sana has their own installation and every store - of which a customer can have multiple - is a separate web application, not connected in any way to other customers that use Sana. Sana records orders that are placed through every Sana online store[1]. In my dataset, each seller is identified with a random, unique identifier and I know only their headquarters' country and industry, due to privacy considerations. In the majority of cases, the currency and the value of the transactions in the original currency are recorded. Importantly, e-commerce sites can support multiple currencies and serve buyers from many countries. However, the only information available regarding the buyers is their currency.

For the purpose of building the dataset for this thesis, transactions of each seller were aggregated on the daily level, per currency, thus the daily number of transactions (for a seller, for a currency) and the daily value of transactions (also for a seller, for a currency), are

[1] This applies only to newer versions of the Sana software (starting in 2019) therefore not every Sana customer is tracked this way as many still use various older versions. Whether the currency and the transaction value are known is also determined by which version of Sana the store runs on.

known. The latter is referred to in e-commerce as the daily Gross Merchandise Value (GMV). The dataset I received contains data for 536 Sana customers and in total about 7,9 million transactions which occurred between the 11th of February 2020 and the 11th of April 2022. The panel is unbalanced, as the period covered differs per seller. For example, it is shorter if they stopped using the Sana Commerce software during the sample period.

I work with a sample of 536 predominantly business-to-business firms that have some notable characteristics. Firms that become customers of Sana likely have higher than average profits if they can afford the investment and higher technological sophistication, especially compared to their peers that do not have an e-commerce system. From a purely technical perspective, they must have a certain, rather modern and costly, type of enterprise resource planning software that can be connected to Sana. For those reasons, I can assume that they tend to be larger than average. This has implications for my results because customers of Sana might be more productive than the average firm in their sector. Therefore they might be more resilient to unfavourable economic conditions. For example, assuming imperfect competition, they could be more capable than their competitors to lower prices in response to a negative demand shock, such as the war in Ukraine. Then, my results would be biased if I were to attempt generalising them to the whole economy. They would imply the change in the daily number of transactions larger than the true population value, because customers of Sana would simply increase their market share, and the change in the daily value of transactions smaller than the population change because their prices would be reduced.

Furthermore, the data I work with covers only transactions placed online. In reality, there exists a great heterogeneity as to what percentage of orders is placed online versus offline between companies. Offline sales include orders placed by phone or in a physical location. The adoption of an online store might also vary over time, and an upward trend is typically observed. In the context of my thesis, this raises the concern that if in the aftermath of the war in Ukraine buyers become systematically more (or less) likely than before to place orders online, it will have an impact on my results.

3.2. Preparation of data

The first key challenge was caused by rows with missing currency information. 228 sellers had missing values for the currency for at least one date. This is a large percentage of the sample and since without knowing the currency I would not be able to convert the transaction values to US dollars, I filled in the majority of the missing currency data points based on two assumptions. First, if for a given seller, some dates have some known currency and other, strictly different dates do not, I can replace the missing currency data points with the known currency. This change affected 136 sellers. Second, if for a given seller the currency is missing for all dates and all these dates are unique[2], I assume that the currency is the official currency in the seller's country, for example, euros for Italy. Only 11 sellers fell under this category. For the remaining 81 sellers that fulfil neither of the above criteria, I dropped all data associated with them.

Leveraging the fact that each online store might process transactions in multiple currencies, I set the seller-currency pair, rather than the seller, to be the cross-sectional unit of observation. For example, if a seller sells in United States dollars and Canadian dollars, I have two separate time series for each variable of that seller. This is desirable because it allows me, albeit imperfectly[3], to group transactions by the country of the buyer as well as the country of the seller. I also dropped 25 seller-currency pairs that had fewer than 30 days of data because they were likely failed e-commerce projects or new online stores with low adoption. In the end, the number of cross-sectional units of observation is 629.

[2] I use the fact that if there were two rows of data for the same seller and the same date, they must be referring to transactions in two different currencies.

[3] In particular, knowing that an online store sells in euros does not precisely determine the country of buyers, I only know that it is in the eurozone. Buyers could also pay in a currency different from that of their home country, although, typically, sellers offer one e-commerce portal per country of distribution.

Another, more technical, challenge arose because the time series for each seller-currency pair in the original dataset contain gaps if no transactions were recorded on a given day. This is very common on weekends due to the business-to-business nature of the transactions but there are also numerous gaps longer than a week when a particular online store processed no transactions. To solve that problem, I added a row with the daily GMV and the daily number of transactions equal to zero for every day, between the first and the last date with data, that was missing. After executing these steps, the dataset had a correct panel structure with day as the time series unit of observation and seller-currency pair as the cross-sectional unit.

I proceeded to adjust the daily GMV values for inflation in the seller country. Because inflation rates change over time and accelerate at the end of my sample period, country-fixed effects cannot be used to sufficiently control for inflation. I used data on the Consumer Prices Indexes from the International Financial Statistics (2022) where the base year is 2010. Due to the unavailability of reliable inflation data, I dropped firms located in Argentina from the sample. Monthly data was available for all other countries except for Australia and New Zealand therefore for those two I used quarterly data. Next, I converted all values of daily GMV to the same currency, the US dollar. To do so, I relied on monthly data that contains the average foreign exchange rate for every USD-local currency pair, available from the Bank of International Settlements (2022). There were a handful of currencies, comprising a small percentage of my dataset, where the value of the exchange rate for April 2022 was not yet available. In these cases, the GMV values from April 2022 were converted based on the exchange rate for March 2022 and this has a negligible impact on the results.

Time series of the daily transaction number and the daily GMV were used to construct past 30-day rolling sums of transactions and GMV for each seller-currency pair. These variables are added to the regressions to control for growth in online store adoption and the wider economy.

Furthermore, I built another supplementary dataset of what I refer to as exposure measures. This thesis uses two, country-level exposure measures to quantify the economic closeness of a country to countries directly involved in the war, namely Ukraine, Russia and Belarus. They are intended to serve as a proxy for how much a country can potentially be impacted by the war. The first exposure measure captures trade relationships and is constructed using aggregate imports and exports data from the World Integrated Trade Solution (2022) database according to the below formula:

$$trade\ exposure\ _{c,\, t=2019} = \frac{value\ of\ imports + exports\ with\ Russia,\ Belarus\ and\ Ukraine_{c,\, t=2019}}{value\ of\ imports + exports\ with\ the\ world_{c,\, t=2019}} * 100\%$$

Where c refers to a country. Data on imports and exports is yearly and for 2019 - before the beginning of the sample period of the Sana dataset and before the Covid-19 pandemic. By construction, the trade exposure measure varies between 0 and 100 which simplifies the interpretation of the results. For example, for Finland the trade exposure measure equals 9.858% which means that in 2019 around 9.9% of Finland's trade was conducted with Russia, Belarus or Ukraine, in relation to Finland's total trade that year. Importantly, encompassing all trade, this measure does not directly correspond to only online transactions. I hypothesise that transactions across borders are less digitised than all business-to-business transactions which would weaken the impact of this exposure measure on my variables of interest.

The second exposure measure is a country's dependence on natural gas from Russia. Its inclusion is motivated by the fact that energy security is very important in the context of the war in Ukraine. While natural gas represents only a fraction of energy resources exported by Russia, it serves as a reasonable proxy in light of data availability constraints. The Russian gas dependence measure is calculated by combining data from the European Statistical Office (2022c and 2022d) and analogously to a more broad energy dependence measure:

$$gas\ dependence_{c,t=2019} = \frac{caloric\ value\ of\ natural\ gas\ imported\ from\ Russia_{c,\ t=2019}}{caloric\ value\ of\ gross\ available\ energy_{c,\ t=2019}} * 100\%$$

Where c indicates the country and all data is for the year 2019. As per the European Statistical Office (2022b), gross available energy is, to put simply, the overall supply of energy available in a country. Both the numerator and the denominator are reported in terajoules which allows for the calculation of the ratio. Interpretation is very similar to the trade exposure measure - for example, the Russian gas dependence of Germany equals 13.915% meaning, roughly, that the total energy demand of Germany in 2019 was covered in almost 14% by imports of natural gas from Russia. The trade exposure and the gas dependence measures have a low correlation equal to 0.204. Their values for all countries are reported in Appendix A.

Because data from Eurostat covers only the European countries, the sample size is considerably reduced when the gas dependence measure is included in regressions. However, running the analysis for a smaller set of countries (excluding the US, the country with the highest representation of firms in my dataset) serves as an additional robustness check.

Using the country-specific exposure measures instead of more precise ones comes with certain limitations. The ideal measure of exposure to the war would be firm-to-firm (seller-to-buyer) specific and capture how much each firm trades with the countries directly involved in the war, at all levels of the supply chain. There undoubtedly exist significant differences between firms in terms of how much they are exposed. For instance, a manufacturer of baked goods that relies heavily on grain from Ukraine or an energy-intensive firm in the chemicals industry are likely exposed more than the average firm. Moreover, I can focus on only one side of the transactions, the buyers, and their exposure to the war.

Nonetheless, constructing a more disaggregated exposure measure was outside the scope of this thesis. A firm-specific measure would be unfeasible due to the lack of available data and even an industry-specific measure could not be employed. While the dataset from Sana Commerce contains information on the seller's industry, their assignment is not consistent with any major standard classification and the categories belong to different levels of classification - some are narrow like "Fashion and apparel", while some are broader such as "Personal and leisure goods". This makes working with the industry data challenging and cumbersome. Therefore, because of time considerations, industry-level exposure measures were not constructed for this thesis.

The last step of preparing the data involves constructing a country-level dataset. To do so, the daily number of transactions and the daily GMV are grouped at the buyer country level and all values are divided by the number of buyers with available data in that country, on a given day. This results in the average daily transaction number and the average daily GMV as additional dependent variables. Due to the use of averages, the changes in larger firms' transaction pattern will be driving the regression results. In the country-level dataset, I also construct equivalents of the past 30-day rolling sums, this time at the country level.

3.3. Description of data

Summary statistics of the key variables from the firm-level and the country-level datasets are reported in tables 3.1 and 3.2. There are 351,266 day-seller-currency observations but only 259,791 of them have non-missing GMV information. However, the GMV data can be reasonably considered missing at random, since this is determined by which version of the Sana software the firm uses which depends on when did they purchase the Sana product.

What should be noted is the high variance of the daily number of transactions and the daily GMV variables. There is a large discrepancy between the most and the least used online stores in the dataset. On average, a Sana store processes 21.38 transactions per day, worth 10,37 thousand 2010 US dollars. In the country-level dataset, the mean of the averaged variables is lower, suggesting that in the countries with fewer sellers, the sellers are less successful. Variances of the averaged variables are lower by construction. The trade exposure and the Russian gas dependence also vary considerably between countries. In the country-level dataset, the trade exposure, which is the only exposure measure used in the relevant part of my econometric analysis, has different moments compared to the firm-level dataset due to the differences in the number of firms representing each country.

	Daily number of transactions	**Log of the daily number of transactions**	**Daily value of transactions**	**Log of the daily value of transactions**	**Trade exposure**	**Russian gas dependence**
Number of observations	351,266	351,266	259,791	259,791	47	21
Mean	21.38	1.55	10.37	1.03	1.66	8.12
Variance	8,342.14	2.51	2,604.67	1.84	1.89	75.32
Minimum	0	0.00	0.00	0.00	0.20	0.00
Maximum	8,279	9.02	7,746.80	8.96	19.65	61.58
Sum	7,508,320	-	2,695,069.13	-	-	-

Table 3.1. Summary statistics for key variables in the firm-level dataset

The daily value of transactions is in thousands of US dollars, adjusted for inflation with the base year 2010. The unit of observation is day and seller-currency pair. Natural logarithms in columns (3) and (5) are calculated for 1 plus the value. Trade exposure describes the percentage of trade with Russia, Belarus and Ukraine for a country, in proportion to its total trade. Russian gas dependence describes the percentage of gross energy needs of a country that is covered by imports of natural gas from Russia and is available only for the European countries. Trade exposure and Russian gas dependence do not vary over time and are constructed with yearly data for 2019 hence their reported number of observations is equal to the number of countries with available data.

To understand the geographical dimension of the data, it is insightful to visualise where the firms in the dataset are located. This is shown in Figure 3.1. where countries are coloured according to the number of Sana Commerce customers that are headquartered there, with countries coloured in grey having none. In my data, this corresponds to the sellers, rather than the seller-currency pairs. 43 countries have at least one Sana customer. An important characteristic is the high concentration of firms in the United States[4] while the second key

[4] Over 20% of sellers in the final dataset (106 out of 517) are located in the US.

region is Western Europe. Notably, countries bordering Russia such as Latvia, Poland and Finland are represented in the sample, as well as some countries very distant from Russia such as Australia and Chile.

	Average daily number of transactions	**Log of the average daily number of transactions**	**Average daily value of transactions**	**Log of the average daily value of transactions**	**Trade exposure**
Number of observations	28,380	28,380	27,581	27,581	47
Mean	15.37	1.90	8.69	1.33	2.27
Variance	939.82	1.89	618.81	1.54	9.52
Minimum	0	0.00	0.00	0.00	0.20
Maximum	706	6.56	1289.33	7.16	19.65

Table 3.2. Summary statistics for key variables in the country-level dataset

The average number of transactions is calculated as the sum of transactions in a given buyer country on a given day, divided by the number of seller-currency pairs for that country with data available for that day. The average daily value of transactions is constructed analogously and reported in thousands of US dollars, adjusted for inflation with the base year 2010. Natural logarithms in columns (3) and (5) are calculated for 1 plus the value. Trade exposure describes the percentage of trade with Russia, Belarus and Ukraine for a country, in proportion to its total trade. It does not vary over time and is constructed with yearly data for 2019.

Lastly, it is worth zooming in on the time series dimension of the data, in particular in the context of countries differentially affected by the war. Figure 3.2. offers that possibility by showing the logarithm of the average daily transaction number for each date between the 1st of January 2022 and the 11th of April 2022 at the country level and displaying two lines depending on the war exposure value of the country. The yellow line shows values averaged across countries with above mean trade exposure[5] and the green line is for the remaining countries. There is a very strong weekly pattern, typical for business-to-business transactions. Transactions most often take place on Mondays, especially in highly-exposed countries. The two time series do not fully overlap suggesting systematic differences between the two groups of countries (and firms). There is also an upward trend for little-exposed countries.

[5] These countries are: Czech Republic, Ecuador, Estonia, Finland, Germany, Hungary, Israel, Italy, Latvia, Morocco, the Netherlands, Poland, Romania and Tunisia.

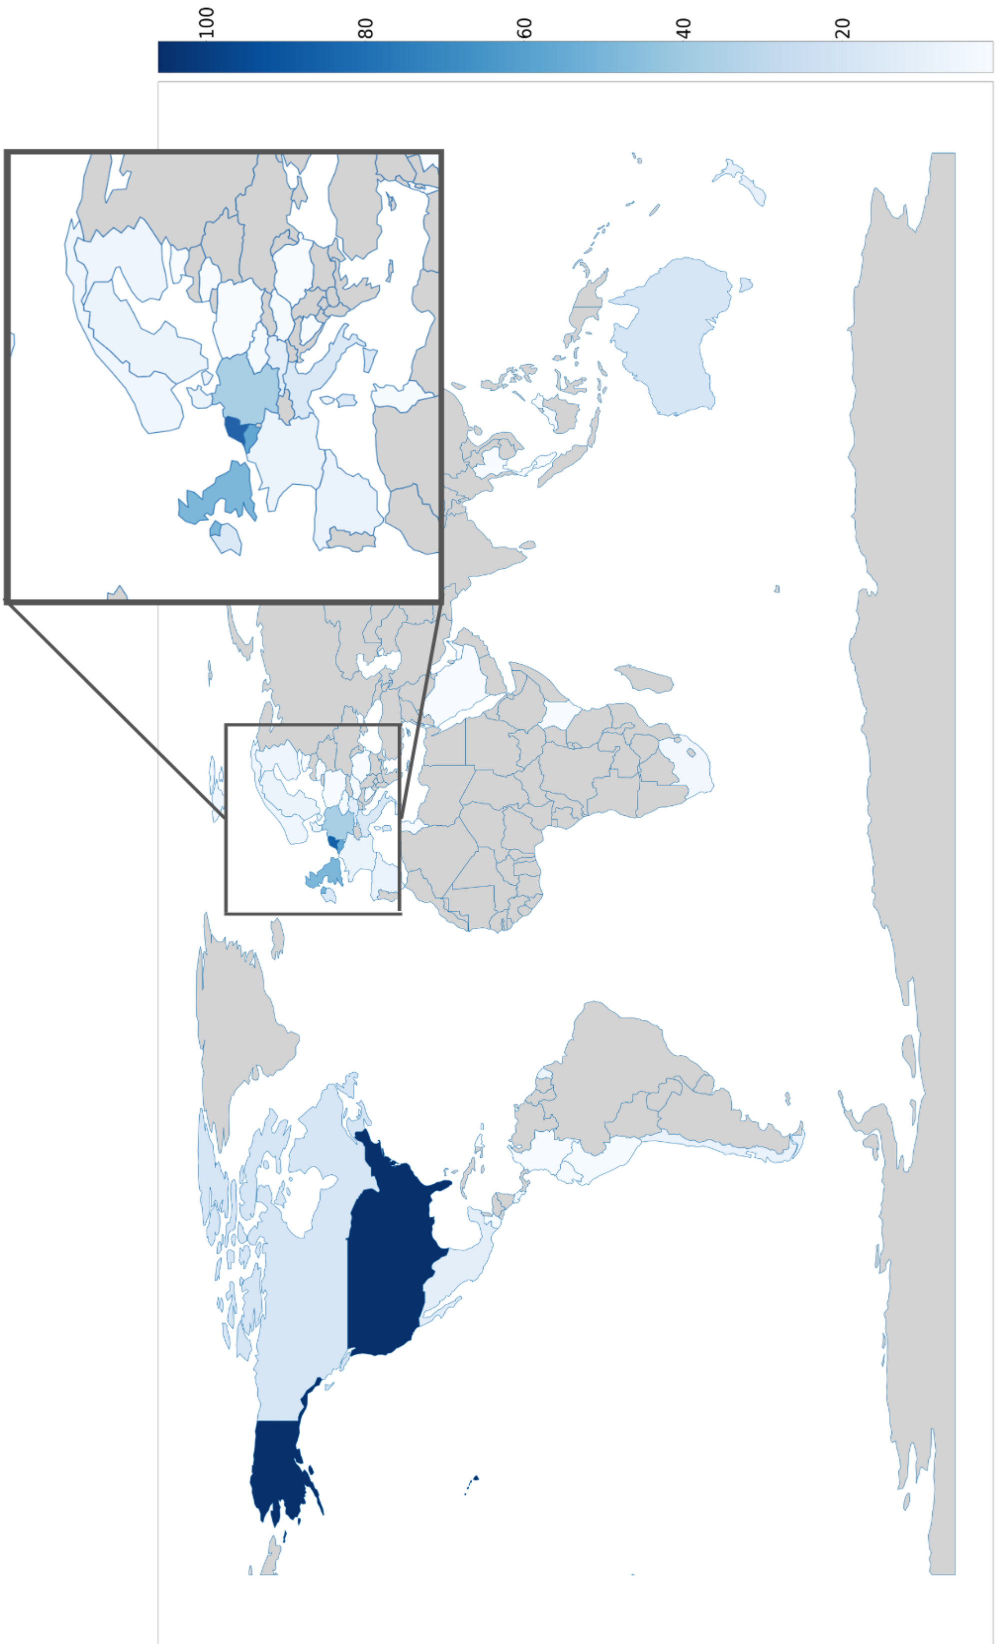

Figure 3.1. The geographical distribution of firms included in the Sana Commerce dataset

Darker shades indicate a higher number of firms in the country. Countries shaded in grey are not represented by any firm in the dataset.

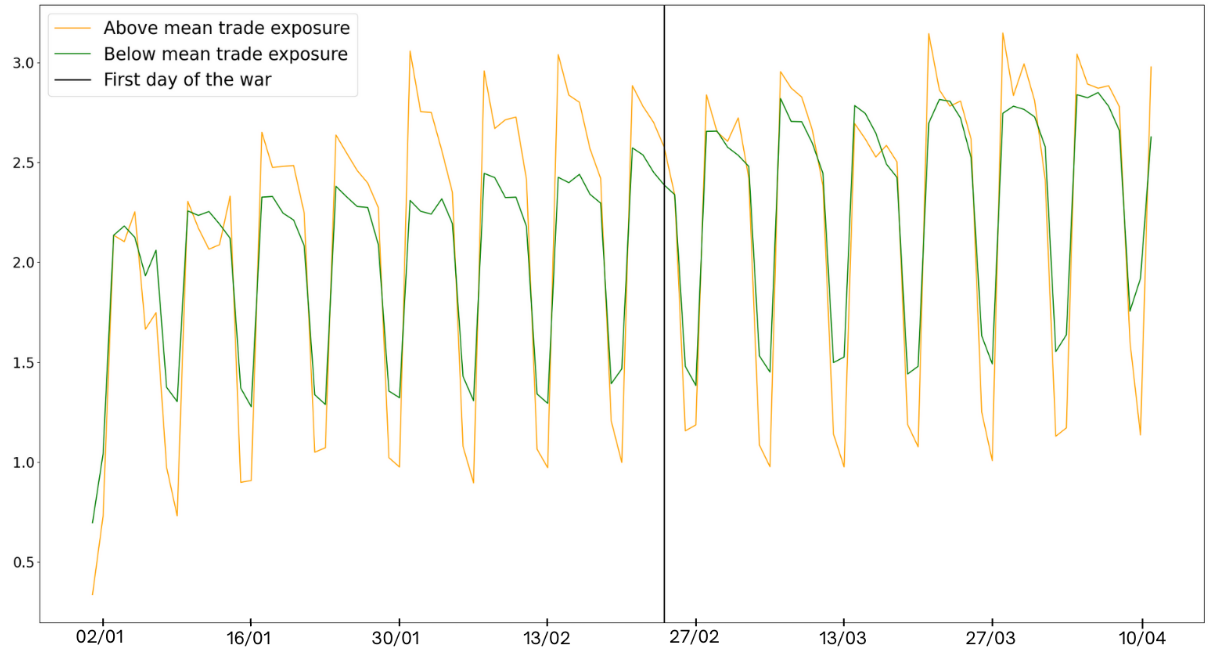

Figure 3.2. Averages of the logarithm of the average daily number of transactions at the country level, for countries with low (yellow line) and high (green line) exposure to the war

Dates on the horizontal axis correspond to every other Sunday. The grey vertical line marks the first day of the full-scale war in Ukraine.

4. Empirical Specifications

4.1. Regression equations

The core of my work involves estimating baseline regressions of the following form, with Y_{fkt} being the logarithm of either the number of transactions or the value of transactions:

$$Y_{fkt} = \beta_0 + \beta_1 war_t + \beta_2 exposure_c + \beta_3 war_t * exposure_c$$
$$+ \beta_4 time\ controls_t + \beta_5 firm\ controls_{fk} + e_{fkt} \quad (1)$$

Where fk denotes a seller-currency pair as described in Chapter 3, t denotes day and c denotes the country of the buyers for the seller-currency pair. The dependent variable is the logarithm of the daily number of transactions for one firm in one currency, which represents the extensive transaction margin, or the logarithm of the daily value of transactions which represents the transaction intensive margin. The choice of the logarithmic form is motivated by the shape of the distributions of the underlying variables which feature frequent outlier values[6]. The dummy variable war equals 1 for dates on and after the 24th of February 2022 which was the first day of the full-scale Russian invasion in Ukraine while the variable $exposure$ captures to what degree the country of the buyers is potentially exposed to the war, as specified in Chapter 3. β_1 and β_3 are the primary coefficients of interest in this thesis with the former describing how much the transaction number or value changed during the war compared to the pre-war period, all else equal. The latter is included to measure how much that change differs depending on the value of the war exposure.

$Time\ controls$ is a vector of variables added to account for the weekly and yearly patterns of transactions. It includes dummy variables for the day of the week and the month of the year. $Firm\ control$s is a vector that includes firm-specific information, namely the industry, and the past 30-day rolling sum of the dependent variable which is seller-currency specific. The rolling sum allows me to account for the growth both at the overall economy level and - more importantly for the period studied - the growth in online stores' adoption.

In addition to the firm-level analysis, I estimate a country-level equivalent of the above equation, maintaining daily data frequency. Corresponding to equation (1) is the following:

$$ln(transaction\ number)_{ct} = \beta_0 + \beta_1 war_t + \beta_2 exposure_c + \beta_3 war_t * exposure_c + \beta_4 time\ controls_t + \beta_5 rolling\ sum_c + e_{ct} \quad (2)$$

Analogous regression is estimated for the country-level transaction value. $rolling\ sum$ is the only country-level control variable because industry fixed effects cannot be included and it is, as before, the sum of the past 30 days of the dependent variable's values.

Lastly, building upon equation (1), I estimate an event study-style regression, for each of the two dependent variables. I do so by including $war\ days$, a vector of 57 date dummies for

[6] This is typical for business-to-business firms as they tend to have fewer customers and fewer orders than business-to-consumer firms while having a high average order value. Only a few transitions might constitute a bulk of a firm's monthly revenue.

days between the 14th of February 2022 and the 11th of April 2022, in place of the single war period dummy variable. Hence, β_1 is now a vector of 57 coefficients capturing the departure from the pre-war transaction number (or value) for individual days of the war:

$$Y_{fkt} = \beta_0 + \beta_1 war\ days_t + \beta_2 exposure_c + \beta_3 war_t * exposure_c + \beta_4 time\ controls_t + \beta_5 firm\ controls_{fk} + e_{flt} \quad (3)$$

4.2. Assumptions and their assessment

All regressions in this thesis are estimated by OLS which relies on the fundamental assumption that the expected value of the error term, conditional on the independent variables, equals 0. If that holds, the OLS estimates are unbiased and consistent. I use robust standard errors due to heteroscedasticity concerns and no clustering as neither of the observed variables would allow me to construct meaningful clusters.

The key reason why OLS assumptions might be violated is the omitted variable bias. It occurs when there exists a variable correlated with both the dependent and the independent variable. The identification strategy of this thesis is built around the idea that the war in Ukraine was an unexpected event and can be seen as an exogenous shock. In such a setting, an omitted variable correlated with the war period dummy is not plausible. However, I must also assume that the country-level measures of exposure to the war and the online transactions are not driven by the same factors, to have trust in the estimated coefficients $\widehat{\beta_2}$ and, crucially, $\widehat{\beta_3}$.

While this might be unrealistic, the concern is mitigated by the fact that I use panel data. Thanks to that, I can control for unobserved characteristics of firms, yet only at the country level. This is done implicitly, by including time-invariant country-level exposure measures which will pick up some of the variation in unobserved factors. However, the inclusion of firm fixed effects is not possible with my research design, leaving a significant portion of the between-firms variation unexplained.

Another concern could stem from potential omitted variables that vary over time which standard panel data methods do not allow to control for. However, I have calculated the exposure measures using data for 2019, before the start of the sample period. Thanks to this solution, I am not concerned about biases introduced by time-varying factors that could impact both the exposure measures and the daily transactions number or value.

Another relevant reason to doubt the regression results is the measurement error. Given the nature of the data I work with, it can pose a serious problem. For example, e-commerce sites can go offline due to a technical problem and process no transactions while buyers are willing to place orders. Luckily this affects only the dependent variables and causes no bias if the error is not correlated with any of the dependent variables. In the context of the above example, as long as technical problems did not become more common during the war in Ukraine, measurement error should not bias the estimated coefficients.

Lastly, a specific limitation of the country-level analysis should be discussed. It comes from the fact that countries are represented by vastly different numbers of observations. The aggregation is performed by grouping by the buyer country, which is determined based on the currency of the transactions, and while there are hundreds of firms in the sample that sell in the US dollar, only a few sell, for example, in the Danish krone. As a result, each seller-currency pair from a country which is represented by a few seller-currency pairs, has a disproportionately higher weight in the country-level regression, compared to firm-level regression. Relatedly, because the dependent variables are averaged, within any given country, larger firms have more impact on the movement of variables over time than smaller firms. Having these issues in mind, it is still valuable to run the country-level analysis to evaluate the robustness of the results.

5. Main Results

5.1. Firm-level results

First, I focus on analysing the pattern of the number of transactions in response to the war in Ukraine. Results of regressions where the dependent variable is the daily number of transactions for a seller-currency pair are reported in Table 5.1. A different set of control variables is included in each column and the baseline regressions' results, as per equation (1) specified in Chapter 4, are reported in columns (4) and (7).

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
War dummy variable	0.250*** (0.015)	0.185*** (0.013)	0.166*** (0.013)	0.175*** (0.013)	0.222*** (0.017)	0.185*** (0.014)	0.166*** (0.015)
Trade exposure	-0.023*** (0.002)	-0.022*** (0.002)	-0.022*** (0.002)	-0.013*** (0.001)			
War dummy variable*trade exposure	-0.014** (0.006)	-0.023*** (0.006)	-0.023*** (0.006)	-0.026*** (0.005)			
Gas dependence					0.022*** (0.000)	0.016*** (0.000)	0.019*** (0.000)
War dummy variable*gas dependence					-0.001 (0.002)	-0.006*** (0.001)	-0.007*** (0.001)
Rolling sum	no	yes	yes	yes	no	yes	yes
Time controls	no	no	yes	yes	no	no	yes
Industry controls	no	no	no	yes	no	no	yes
Number of observations	351,266	333,025	333,025	333,025	218,803	207,725	207,725

Table 5.1. Estimates obtained from regressions with the daily number of transactions as the dependent variable

The dependent variable is the logarithm of the daily number of transactions for each seller-currency pair. In columns (1)-(4) trade exposure is included as the measure of the exposure to the war. It captures the percentage of trade with Russia, Belarus and Ukraine for a country, in proportion to its total trade. In columns (5)-(7) the Russian gas dependence is included instead. It measures the percentage of gross energy needs of a country that is covered by imports of natural gas from Russia and is available only for the European countries. War dummy variable equals 1 for all dates on and after the first day of the war in Ukraine, 24-02-2022. Rolling sum is a past 30-day rolling sum of the number of transactions for a given seller-currency pair. Time controls is a vector of 6 dummy variables for the days of the week and 11 dummy variables for the month of the year. Industry controls is a vector of 18 dummy variables for industries.
Robust standard errors are reported in parenthesis. Coefficients significant at 10%, 5% and 1% are marked with *, ** and ***, respectively. Results are rounded to 3 decimal places.

The full results of the regression in column (4) are presented in Appendix B. The estimated coefficients for the dummy variable which equals 1 for dates during the war in Ukraine and 0 otherwise (first row) are remarkably consistent and highly statistically significant for both measures of exposure to the war. The results suggest that the war period is, on average, associated with an increase of between 17% to 25% in the daily number of transactions. In absolute terms that is a highly economically significant increase from, on average, 21.38 transactions per day to between 25.01 and 26.73 transactions per day.

Notably, for the trade exposure as well as gas dependence, the inclusion of additional controls tends to reduce the magnitude of the coefficient associated with the war period. This suggests that especially the results in columns (1) and (5) are affected by omitted variable bias where the bias has a positive sign i.e. overestimation. This is consistent with my hypothesis - the inclusion of the rolling sum of past transactions, which helps to control for the general upward trend in the number of transactions, is necessary if the number of past transactions is positively correlated with their present daily sum and with the war period variable.

Furthermore, there exist robust but not very sizable correlations between the trade exposure and gas dependence measures and the daily number of transactions for the period before the war. Firms in countries highly reliant on trade with Russia, Belarus and Ukraine tend to have fewer transactions on average. One percentage point increase in the trade exposure is associated with a 1.3% decrease in the daily number of transactions, as reported in column (4). For an increase of one standard deviation in trade exposure, the reduction in the daily number of transactions equals approximately 1.8%. On the other hand, for gas dependence, one standard deviation increase is associated with an increase of about 16.5% in the number of transactions. This can likely be explained by the fact that firms from different industries in the sample are represented in different countries in a non-random way.

Lastly, it is critical to discuss how the degree of war exposure differentially affects the change in the number of transactions after the start of the war. My results suggest that the higher the war exposure, the smaller the increase in the number of transactions described above. Importantly, I reach this conclusion for both exposure measures. In the case of a one standard deviation increase in the trade exposure, the estimated coefficient for the interaction term implies a 3.57% smaller change associated with the war period, while for a one standard deviation increase in the Russian gas dependence, that value equals approximately 6.08%. When these estimates are related back to the absolute numbers, one can notice that the values are economically not very significant, but they are statistically robust.

Next, I describe the main results obtained for the pattern of the daily value of transactions, or the GMV, before I proceed to jointly elaborate on the results in the context of the theoretical framework.

Crucially, my analysis points to a very similar movement on the intensive transactions margin compared to the extensive transactions margin in response to the war, even though the regressions with the GMV are estimated on a significantly (around 30%) smaller sample. Relevant estimates for regressions with daily GMV as the dependent variable are reported in Table 5.2. The baseline regressions' results are shown in columns (4) and (7) for the trade exposure and Russian gas dependence, respectively.

The period during the war in Ukraine is associated with a highly statistically significant increase in the daily value of transactions of between 15% to 28%. These values are substantial in economic terms suggesting that the daily GMV was around 2.1 to 3.4 thousand USD higher for the average seller-currency pair when expressed in today's purchasing power.

One difference with the results for the daily number of transactions is that I do not obtain statistically or economically significant differences in the pre-war value of transactions depending on the value of trade exposure. I do, however, find evidence for the differences in transaction value depending on the degree of gas dependence. Analogous to previous results, a higher gas dependence is strongly associated with a higher daily value of transactions.

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
War dummy variable	0.280*** (0.015)	0.159*** (0.013)	0.156*** (0.013)	0.156*** (0.013)	0.213*** (0.016)	0.155*** (0.014)	0.149*** (0.015)
Trade exposure	-0.002 (0.002)	-0.002 (0.001)	-0.002 (0.001)	-0.000 (0.001)			
War dummy variable*trade exposure	-0.019*** (0.005)	-0.016** (0.005)	-0.016** (0.005)	-0.017*** (0.005)			
Gas dependence					0.010*** (0.000)	0.004*** (0.000)	0.007*** (0.000)
War dummy variable*gas dependence					0.000 (0.001)	-0.004** (0.001)	-0.005*** (0.001)
Rolling sum	no	yes	yes	yes	no	yes	yes
Time controls	no	no	yes	yes	no	no	yes
Industry controls	no	no	no	yes	no	no	yes
Number of observations	259,791	244,921	244,921	244,921	161,067	152,075	152,075

Table 5.2. Estimates obtained from regressions with the daily GMV as the dependent variable

The dependent variable is the logarithm of the daily value of transactions (GMV) for each seller-currency pair. In columns (1)-(4) trade exposure is included as the measure of the exposure to the war. It captures the percentage of trade with Russia, Belarus and Ukraine for a country, in proportion to its total trade. In columns (5)-(7) the Russian gas dependence is included instead. It measures the percentage of gross energy needs of a country that is covered by imports of natural gas from Russia and is available only for the European countries. War dummy variable equals 1 for all dates on and after the first day of the war in Ukraine, 24-02-2022. Rolling sum is a past 30-day rolling sum of the value of transactions for a given seller-currency pair. Time controls is a vector of 6 dummy variables for the days of the week and 11 dummy variables for the month of the year. Industry controls is a vector of 18 dummy variables for industries.
Robust standard errors are reported in parenthesis. Coefficients significant at 10%, 5% and 1% are marked with *, ** and ***, respectively. Results are rounded to 3 decimal places.

Lastly, the sign of interaction term between the war date dummy and both measures of the war exposure is consistently negative. In other words, firms in highly-exposed countries experience a smaller change in the daily value of transactions compared to firms in little-exposed countries. To quantify, my results indicate that an increase of one standard deviation in the trade exposure measure is, on average, associated with a 2.34% smaller increase in the daily value of transactions during the war period. Meanwhile, an increase of one standard deviation in the Russian gas dependence coincides with, on average, a 4.34% smaller change in the daily value of transactions. The values obtained this way are remarkably close to analogous values in the case of the daily number of transactions which, as reported above, equal 3.57% for trade exposure and 6.08% for gas dependence.

Results presented in this section suggest that after the full-scale Russian invasion of Ukraine, the average daily number, as well as the average daily value, of online business-to-business transactions increased quite substantially, at least in the short term[7]. Even more strikingly, I find evidence that these increases were slightly less pronounced in countries that have strong economic ties with the countries involved in the war.

Both of those central conclusions are not expected considering the theoretical framework outlined in Chapter 2. While potential explanations are presented in detail in the following chapter, it is worth pointing to the high degree of similarity between regressions that include different exposure measures which might suggest the robustness of the results. Regressions that include the gas dependence are run on a restricted sample of only buyer countries located in Europe, yet key results obtained are close to the results which rely on the trade exposure.

5.2. Country-level results

In this section, I repeat parts of the firm-level analysis using a modified dataset where the dependent variables are aggregated to the country level and averaged to account for differences in the number of observations for each country. As a result of the aggregation, it is impossible to include industry fixed effects and the sample size decreases dramatically. For brevity, I include only one type of measure of the war exposure, namely the trade exposure.

[7] The sample period ends on the 11th of April 2022, less than two months after the start of the war therefore nothing can be concluded about the medium- and long-run impact of the war on online business-to-business sales.

Estimates from regressions with the average of the logarithm of the daily transaction number as the dependent variable are presented in Table 5.3., while analogous results for the average of the logarithm of the daily transaction value are presented in Table 5.4.

	(1)	(2)	(3)
War dummy variable	0.438*** (0.039)	0.174*** (0.031)	0.141*** (0.030)
Trade exposure	-0.064*** (0.002)	-0.032*** (0.002)	-0.033*** (0.002)
War dummy variable *trade exposure	0.004 (0.007)	0.002 (0.007)	0.002 (0.005)
Rolling sum	no	yes	yes
Time controls	no	no	yes
Number of observations	28,380	27,046	27,046

Table 5.3. Estimates obtained from regressions with the average number of transactions as the dependent variable

The dependent variable is the average of the logarithm of the daily transaction number for each buyer country. Trade exposure captures the percentage of trade with Russia, Belarus and Ukraine for a country, in proportion to its total trade. War dummy variable equals 1 for all dates on and after the first day of the war in Ukraine, 24-02-2022. Rolling sum is a past 30-day rolling sum of the number of transactions for a given buyer country. Time controls is a vector of 6 dummy variables for the days of the week and 11 dummy variables for the month of the year.
Robust standard errors are reported in parenthesis. Coefficients significant at 10%, 5% and 1% are marked with *, ** and ***, respectively. Results are rounded to 3 decimal places.

The first key insight is that the estimates for dates during the war are remarkably close to their equivalents from the firm-level analysis. The results imply that, after the war began, the average daily number of transactions per seller-currency increased by approximately 14.1%, while the average daily value of transactions increased by approximately 18.6%, as reported in column (3) of Tables 5.3. and 5.4. This similarity could imply that the size of the adjustment after the war was comparable across countries - which is explored in Chapter 7 - and across firms of different sizes. This would be consistent with the earlier finding that the degrees of trade exposure and gas dependence do not have a large impact on the size of the war period changes at the firm level. One can also notice the same as before pattern of the coefficients decreasing as more control variables are added.

	(1)	**(2)**	**(3)**
War dummy variable	0.600*** (0.041)	0.158*** (0.033)	0.186*** (0.031)
Trade exposure	-0.035*** (0.002)	-0.016*** (0.002)	-0.017*** (0.002)
War dummy variable *trade exposure	-0.025*** (0.007)	0.008 (0.006)	0.008* (0.005)
Rolling sum	no	yes	yes
Time controls	no	no	yes
Number of observations	27,581	26,276	26,276

Table 5.4. Estimates obtained from regressions with the average value of transactions as the dependent variable

The dependent variable is the average of the logarithm of the daily transaction value for each buyer country. Trade exposure captures the percentage of trade with Russia, Belarus and Ukraine for a country, in proportion to its total trade. War dummy variable equals 1 for all dates on and after the first day of the war in Ukraine, 24-02-2022. Rolling sum is a past 30-day rolling sum of the value of transactions for a given seller-currency pair. Time controls is a vector of 6 dummy variables for days of the week and 11 dummy variables for month of the year.
Robust standard errors are reported in parenthesis. Coefficients significant at 10%, 5% and 1% are marked with *, ** and ***, respectively. Results are rounded to 3 decimal places.

I, again, find statistically significant evidence that firms with a higher degree of trade exposure had a lower average daily number of transactions. A one standard deviation increase in the trade exposure is associated with a 10.18% decrease in the average daily number of transactions. At the same time, I find that the average daily value of transactions was also lower in the highly-exposed countries - for one standard deviation increase in trade exposure, the drop in average daily GMV is around 5.25%. Together, these results imply that the highly-exposed countries are characterised by a higher average order value. This could be counterintuitive in light of the fact that countries in my sample that trade a lot with Russia, Belarus and Ukraine are likely poorer than countries which trade little but can be well explained if we consider the heterogeneity of industries represented in each country.

Finally, the coefficients on the interaction terms between the war period dummy and the trade dependence measure, do not provide conclusive evidence regarding the differential behaviour of either dependent variable for varying degrees of war exposure. This might be indicative of the shortcomings of using the country-level approach which were discussed above - by collapsing all firm-level variation for dozens of firms in certain countries into a single, averaged series and comparing it to countries where only a few firms operate, it could be impossible to recover a statistically significant effect.

5.3. Event study results

In this section, I present the results of the event study style regressions that are based on the approach implemented by Dunn et al (2020). I use the firm-level dataset to run regressions of two types, each for the two main dependent variables. In the first type, only a set of 57 dummy variables for every available date during the war is included, alongside controls. The key results are graphically summarised in Figure 5.1. In the second type of regressions, I include the same date dummies as well as interactions between them and the trade exposure measure. I report the coefficients of the interaction terms obtained this way in Figure 5.2.

Focusing on the date dummy variables alone, the key conclusion is that the estimated coefficients are generally above zero throughout the whole period studied. This is not surprising and in line with my previous results. The coefficients obtained from regressions with the daily number of transactions and the daily value of transactions as dependent variables are extremely similar - the correlation between the two series is equal to 0.958. Interestingly, the weekly pattern of data prevails even though control variables for the days of the week are included. This might suggest that weekly fluctuations became stronger during the war in Ukraine. Moreover, one can easily notice an upward trend which accelerates in the second half of March. At the beginning of April, the increase in the daily number and value of transactions exceeds 40%.

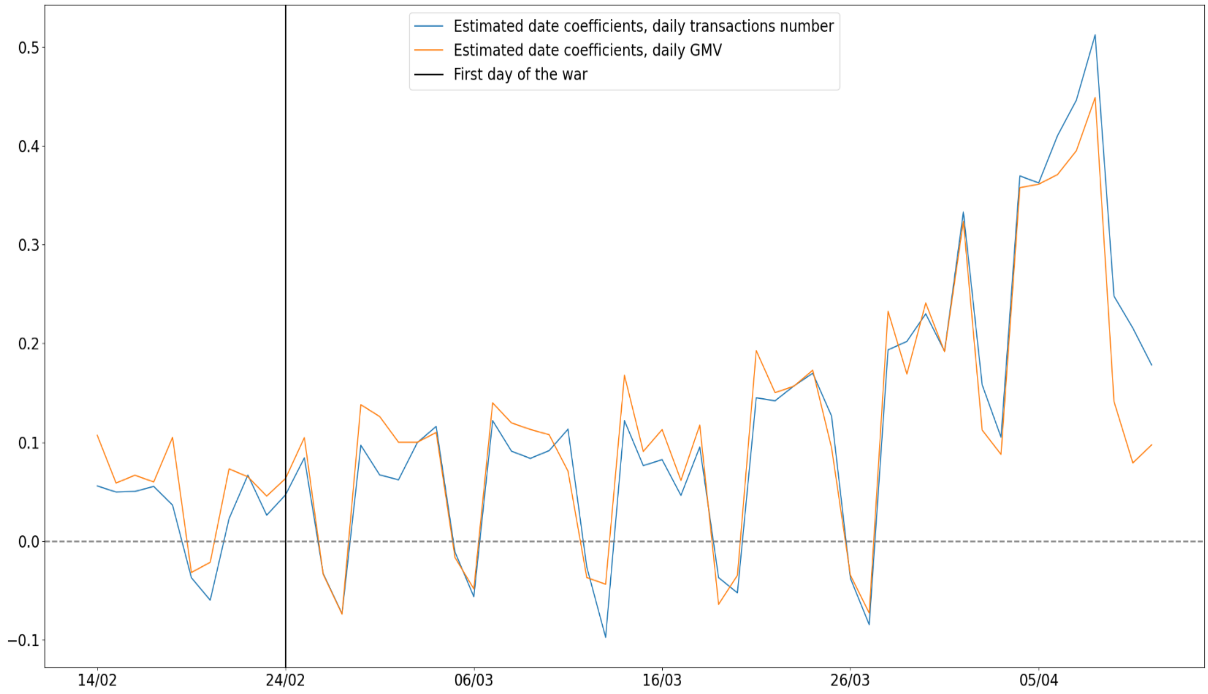

Figure 5.1. Estimated coefficients for each of the dummies for dates during the war

The vertical axis gives estimated values of coefficients and the horizontal axis shows dates associated with each dummy variable. The blue line traces the coefficients obtained from regressions where the logarithm of the daily number of transactions, at the firm level, is the dependent variable, while the orange line traces coefficients from an analogous regression for the logarithm of the daily value of transactions is the dependent variable. Both regressions include controls for the rolling sum of transactions or GMV in the past 30 days, the day of the week and the month of the year, as well as for industry.

What must be noted is that the coefficients are estimated at above 0 for most dates before the beginning of the war. While a detailed analysis of when the effect of the war became detectable is beyond the scope of this thesis, it might suggest that the impact of the war on the economy had started before the invasion took place due to rational expectations of economic agents. However, it might also suggest that there was another force driving transactions' number and value upward, on top of the effects of the war.

When it comes to the interaction terms' coefficients, their pattern is also similar between the transaction number and transaction value regressions with a correlation equal to 0.846. Unsurprisingly, almost all estimated coefficients are below 0 with one notable outlier on the 11th of April 2022. Both series presented in Figure 5.2. exhibit rather high variance which cannot be explained by the weekly pattern of business-to-business transactions, suggesting that the differential reaction to the war depending on the trade exposure measure is not very strong or consistent. All in all, the event study analyses support my previous conclusions and add valuable depth to the results.

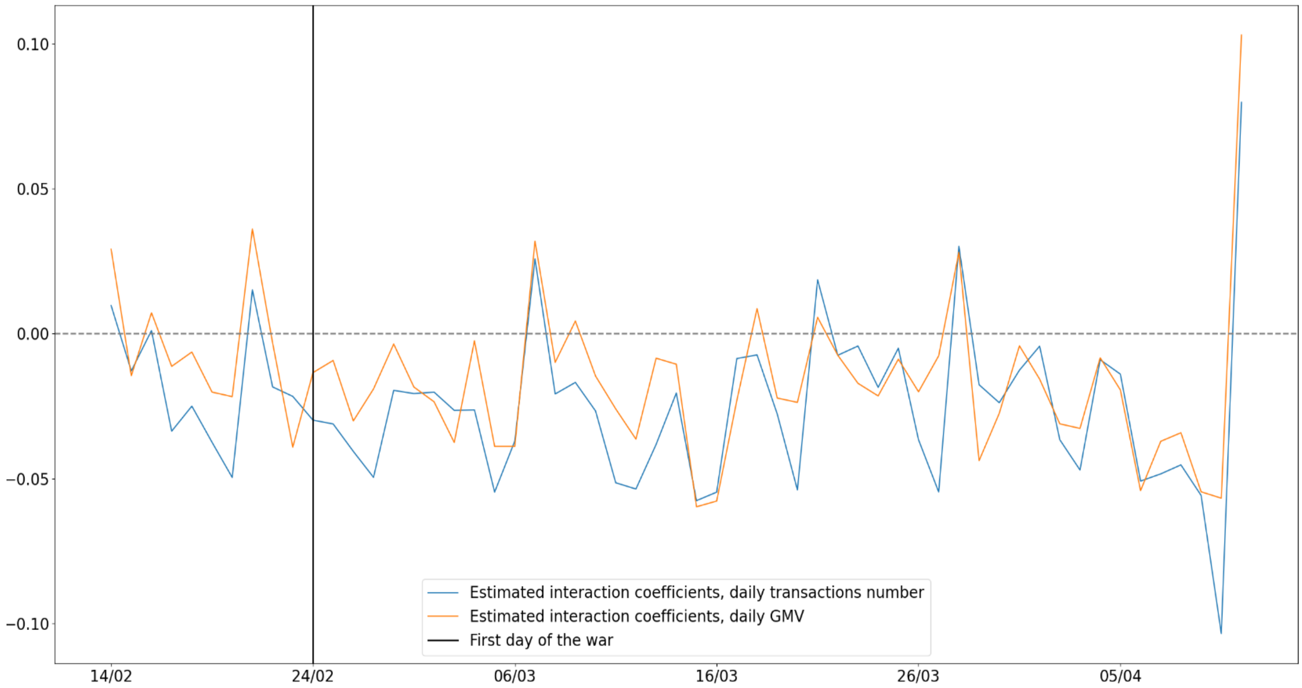

Figure 5.2. Estimated coefficients for interaction terms between trade exposure and each of the dummies for dates during the war

The vertical axis gives estimated values of coefficients and the horizontal axis shows dates associated with each dummy variable. The blue line traces the coefficients obtained from regressions where the logarithm of the daily number of transactions, at the firm level, is the dependent variable, while the orange line traces coefficients from an analogous regression for the logarithm of the daily value of transactions is the dependent variable. Trade exposure is the percentage of trade with Russia, Belarus and Ukraine for a country, in proportion to its total trade. Both regressions include controls for the rolling sum of transactions or GMV in the past 30 days, the day of the week and the month of the year, as well as for industry.

6. Mechanisms

In this chapter, I briefly describe the most plausible potential explanations that might help with understanding the results presented above. This is important considering that my conclusions are not aligned with ex-ante expectations. Yet, there could be a number of reasonable contributing factors which stems from the complexity of the questions I tackle and the novelty of my research approach.

First and foremost, the results I obtain possibly do not hold for the whole economy due to the unique sample of firms I work with. The sample includes mostly business-to-business firms, medium or large and the industries they operate in do not cover the whole range of major economic activities. Moreover, as highlighted in Chapter 3, firms that purchase products of Sana Commerce are on average more technology-intensive and likely more profitable than other firms. Further to this point, it is unclear how closely online business-to-business sales follow the path of the aggregate demand. One cannot preclude that what my results reflect is a substitution effect between online and offline transactions which could occur for reasons completely unrelated to the war in Ukraine, for example, due to increases in the number of coronavirus infections. All in all, the conclusions of this thesis cannot be generalised without imposing strong and unrealistic assumptions.

Relatedly, with my research design, it is not entirely possible to isolate the impact of the war in Ukraine from the impact of other events affecting the economy. As my event analyses show, the highest coefficients of the day dummy variables are reported for early April. This implies that the statistically significant positive coefficients on the single war period dummy reported in Sections 5.1 and 5.2 are driven by the large April increases in the daily transaction number and value, already more than a month after the war started. This could be interpreted as evidence that the results should be attributed to another event. In particular, rising inflation expectations and uncertainty, which the war exacerbated but did not cause, could partially explain increased demand, even leading to so-called ‘panic buying’. For example, Coibion, Gorodnichenko, and Ropele (2019) give compelling evidence about how firms’ decisions are affected by increased inflation expectations. Firms reduce employment and capital which could be consistent with increasing the demand for inputs and therefore with my results.

Alternatively, the increased demand that I observe for April could be understood as a ‘sign of relief’ of business-to-business buyers. Late March was arguably the time when it became apparent that scenarios of a nuclear war or wider conflict were averted[8]. As news articles document, the public debate shifted to consider the possibility that Russia will be defeated (NBC News, 2022, March 30; The Washington Post, 2022, March 30). Such explanations highlight the fact that one should always keep in mind what the appropriate counterfactual is when trying to analyse the impact of a significant one-off event such as the war in Ukraine. This is the more relevant, the longer the period in question. Even though my results point to a correlation between the war exposure and the changes in the transaction number and value, the war exposures themselves are not assigned at random and they might be related to other relevant factors. This would amount to an omitted variable bias in the regressions.

The last point that deserves attention is whether it is convincing that the war impacted online business-to-business transactions more strongly in less exposed countries. A potential issue is the crude choice of a single dummy variable that applies to dates after the start of the war and the first date dummy variable used in the event analyses, the 14th of February 2022.

Due to limits on the scope of this thesis, I do not test for the possibility that the adjustment of the number and value of transactions started earlier in the most exposed countries compared to the least exposed countries. In other words, the anticipated war could be affecting the highly exposed firms well before the first bombs were dropped on Kyiv, therefore its effect cannot be fully captured using the 24th of February 2022 cut-off date. In that case, my results can suggest incorrectly that the increase in demand was weaker in countries with strong economic ties with Russia, Belarus or Ukraine. This could be plausible considering that US President Joe Biden warned of the planned Russian invasion as early as December 2021 (The Associated Press, 2021, December 04), under the assumption that these warnings were perceived as more credible in the highly-exposed countries.

---

[8] By the 3rd of April 2022 Russian troops retreated from Kyiv Oblast which was perceived as a major sign that Ukraine could be able to withstand the invasion (Reuters, 2022, April 06).

7. Robustness Checks

7.1. Alternative levels of aggregation along time series dimension

For the first set of robustness checks, I aggregate my firm-level dataset to weekly and monthly levels to replicate the analysis from Table 5.1, columns (1) and (4). Accordingly, the dependent variable in both cases is the (weekly or monthly) number of transactions.

The dataset with weekly aggregation ends on the 10th of April 2022, one day earlier than the main dataset. The monthly dataset ends on the 31st of March 2022 which must be kept in mind when interpreting the results. Due to technical limitations, dummy variables for the dates during the war also capture different time windows - from Monday, the 21st of February 2022 in the weekly dataset and from the 1st of February 2022 in the monthly dataset. The rolling sum variables measure the sum of transactions from the past four weeks in the weekly dataset and the past two months in the monthly dataset, for each seller-currency pair.

	(1)	(2)
War dummy variable	0.424*** (0.048)	0.307*** (0.044)
Trade exposure	-0.003 (0.006)	0.009 (0.006)
War dummy variable *trade exposure	-0.028 (0.020)	-0.040** (0.018)
Rolling sum	no	yes
Time controls	no	yes
Industry controls	no	yes
Number of observations	50,401	47,885

Table 7.1. Estimates obtained from regressions on data aggregated weekly with the weekly number of transactions as the dependent variable

The dependent variable is the logarithm of the weekly number of transactions for each seller-currency pair. Trade exposure captures the percentage of trade with Russia, Belarus and Ukraine for a country, in proportion to its total trade. War dummy variable equals 1 for weeks which last day (Sunday) occurred after the first day of the war in Ukraine, 24-02-2022. Rolling sum is a past 4-week rolling sum of the number of transactions. Time controls is a vector of 11 dummy variables for the month of the year. Industry controls is a vector of 18 dummy variables for industries.
Robust standard errors are reported in parenthesis. Coefficients significant at 10%, 5% and 1% are marked with *, ** and ***, respectively. Results are rounded to 3 decimal places.

Key results obtained using the weekly dataset, which are reported in Table 7.1., are consistent with the main results. Dates during the war in Ukraine are associated with a considerably higher weekly number of transactions and the relationship is statistically highly significant. The coefficient on the trade exposure measure lost the statistical significance, compared to earlier results and its sign became unstable, possibly due to a greatly reduced sample size. At the same time, the coefficient on the interaction term between the trade exposure and date during the war remains robust and negative.

	(1)	**(2)**
War dummy variable	0.339*** (0.102)	0.174* (0.108)
Trade exposure	0.008 (0.014)	0.024* (0.014)
War dummy variable *trade exposure	-0.028 (0.042)	-0.033 (0.038)
Rolling sum	no	yes
Time controls	no	yes
Industry controls	no	yes
Number of observations	11,709	10,452

Table 7.2. Estimates obtained from regressions on data aggregated monthly with the monthly number of transactions as the dependent variable

The dependent variable is the logarithm of the monthly number of transactions for each seller-currency pair. Trade exposure captures the percentage of trade with Russia, Belarus and Ukraine for a country, in proportion to its total trade. War dummy variable equals 1 for months which last day occurred after the first day of the war in Ukraine, 24-02-2022. Rolling sum is the past 2-month rolling sum of the number of transactions. Time controls is a vector of 11 dummy variables for the month of the year. Industry controls is a vector of 18 dummy variables for industries.
Robust standard errors are reported in parenthesis. Coefficients significant at 10%, 5% and 1% are marked with *, ** and ***, respectively. Results are rounded to 3 decimal places.

Results obtained with the monthly dataset, as reported in Table 7.2., also support the conclusion that the war period coincided with a higher number of transactions. The panel is significantly shorter, with a maximum of 26 time series observations, hence the statistical significance of the results drops substantially. Yet, the estimated coefficients on the dummy variable for dates during the war and its interaction with the trade exposure measure have the same size and magnitude as in the regression run with daily data.

7.2. Individual country results

Another robustness check involves running a variant of the baseline regression separately for each buyer country in the sample, given that there is enough data available for that country and the variable of interest - the war period dummy variable - is not omitted. The dependent variable is the logarithm of the daily number of transactions because using the daily value of transactions would allow me to run the regression only for a smaller set of countries due to data limitations. None of the war exposure measures is included as they do not vary across observations for one country. Most regressions preserve the panel structure unless there is only one seller-currency pair for a given country.

The main results of the individual country analysis are depicted in Figure 7.1. where each country is coloured according to the value of the estimated coefficients for the dummy variable associated with dates during the war in Ukraine. Consistent with the earlier results, for most countries that coefficient is positive, but a geographical pattern is hard to discern. Countries with the highest estimated coefficients - coloured in dark blue - include Italy, Mexico and South Africa. However, one should keep in mind that different industries are represented in each buyer country. Even though industry fixed effects are included in each regression, in this setting they cannot control for differences in industry representation between countries. The results also reveal that for a number of countries the war period was associated with fewer transactions per day. Notably, those include Romania, Sweden and Latvia which, intuitively, should be highly exposed to the war in Ukraine. More detailed results of the individual country regressions are reported in Appendix C. Importantly, it shows that positive coefficients are more often statistically significant than negative ones.

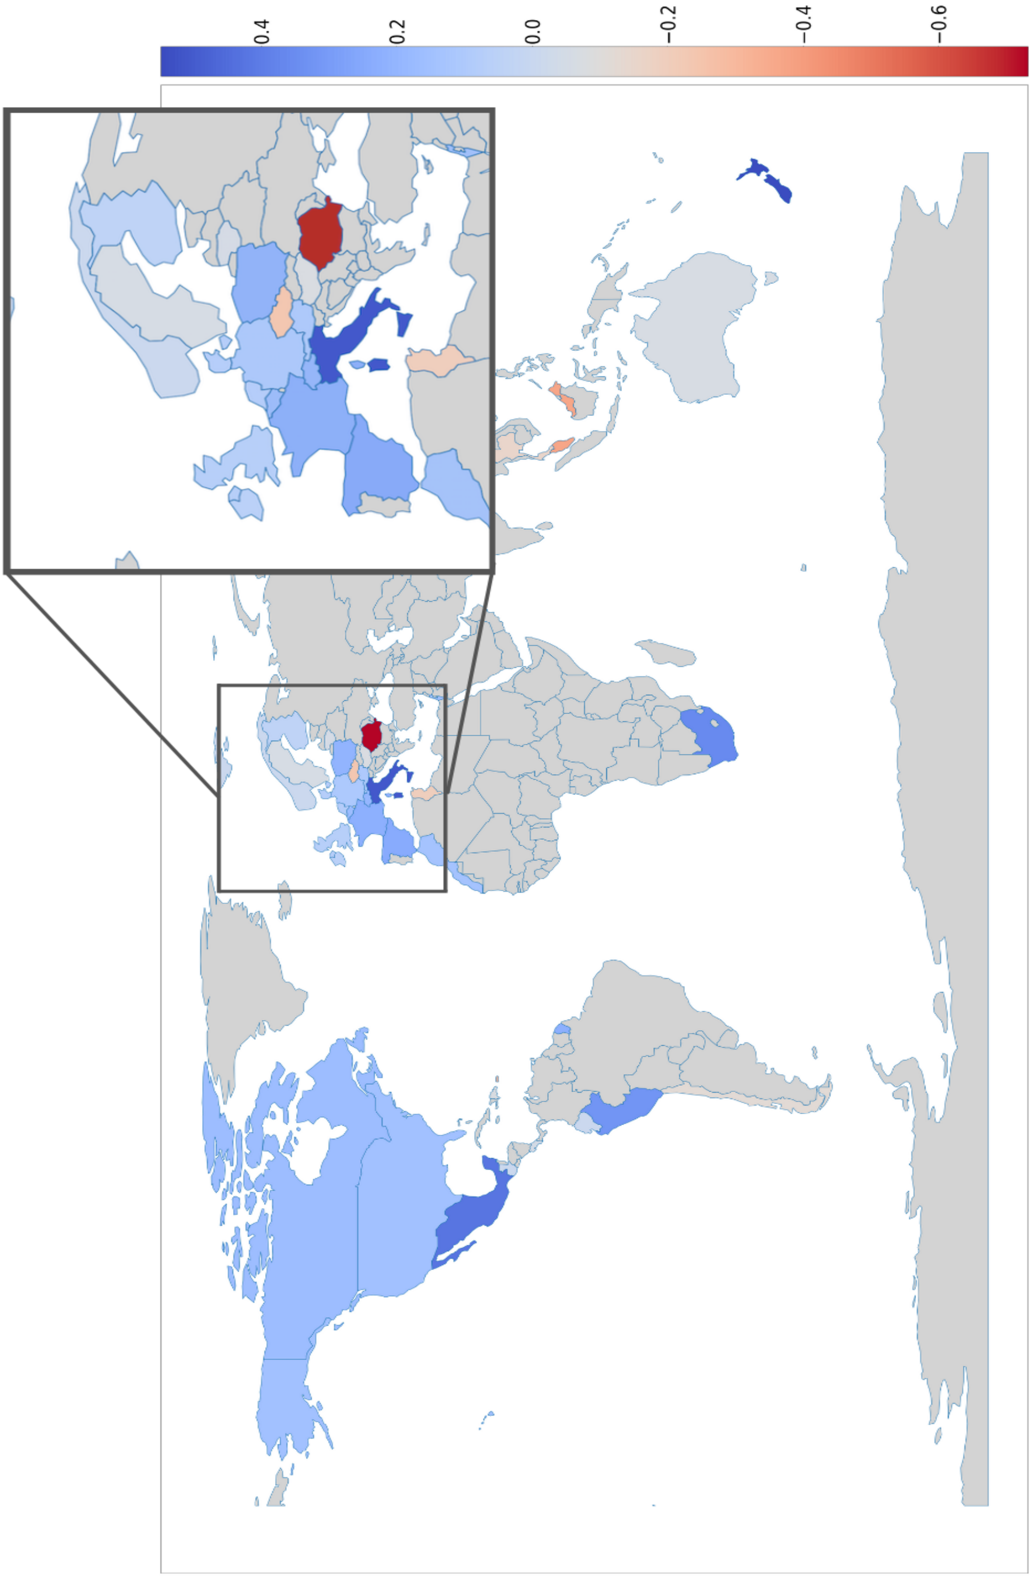

Figure 7.1. Values of the estimated coefficients for the dummy variable representing dates during the war obtained from individual-country regressions with the daily number of transactions as the dependent variable

For countries shaded in blue, the value of the coefficient is positive, while for countries shared orange and red it is negative. For countries shaded in grey, regressions were not estimated due to the lack of available data.

8. Appendix

8.1. Appendix A: Values of exposure measures for each country

Country	Trade exposure	Gas dependence	Country	Trade exposure	Gas dependence	Country	Trade exposure	Gas dependence
Australia	0.200	-	Guatemala	0.454	-	Poland	7.039	8.286
Austria	1.539	0.000	Hungary	4.547	61.581	Romania	3.724	2.636
Belgium	1.786	2.793	Ireland	0.496	0.000	S. Arabia	0.556	-
Bosnia and Herzegov.	2.126	2.889	Israel	2.903	-	Singapore	0.737	-
Canada	0.246	-	Italy	3.017	19.255	S. Africa	0.577	-
Chile	0.663	-	Japan	1.599	-	Spain	1.242	2.285
Colombia	0.484	-	Kenya	1.960	-	Sweden	2.220	0.000
Costa Rica	0.671	-	Latvia	19.645	13.813	Switzerland	0.754	-
Czechia	3.310	20.273	Luxembourg	0.851	4.554	Thailand	0.796	-
Denmark	1.424	0.000	Malaysia	0.531	-	Tunisia	2.698	-
Dominican Republic	0.264	-	Mexico	0.209	-	United Kingdom	1.652	1.601
Ecuador	2.869	-	Morocco	2.638	-	United States	0.824	-
Estonia	10.649	8.854	Netherlands	2.611	17.121	UAE	0.705	-
Finland	9.858	6.679	New Zealand	0.823	-	Uruguay	1.319	-
France	1.532	4.246	Norway	1.519	0.000	Puerto Rico	0.824	-
Germany	2.446	13.915	Peru	0.698	-			

Table 8.1. Full details on the values of the trade exposure and Russian gas dependence exposure measures

Trade exposure is the percentage of trade with Russia, Belarus and Ukraine for a country, in proportion to its total trade. Gas dependence is the percentage of gross energy needs of a country that is covered by imports of natural gas from Russia and is available only for the European countries. Missing data is indicated with "-". Both measures do not vary over time and are calculated based on yearly data for 2019. The countries listed are all countries of buyers. There are more countries of buyers than countries of sellers (customers of Sana).

8.2. Appendix B: Detailed results of baseline regression

War dummy variable	0.175*** (0.013)	February dummy	0.0484*** (0.011)	December dummy	-0.177*** (0.011)	Household goods d.	0.418*** (0.010)
Trade exposure	-0.013*** (0.001)	March dummy	0.042*** (0.011)	Agriculture dummy	0.121*** (0.013)	Leisure goods d.	0.496*** (0.008)
War dummy var.*trade exposure	-0.026*** (0.005)	April dummy	0.007 (0.011)	Automotive parts dummy	0.140*** (0.012)	Machinery dummy	0.104*** (0.008)
Rolling sum	0.000*** (0.000)	May dummy	0.022** (0.011)	Business services d.	-0.537*** (0.025)	Media dummy	0.740*** (0.033)
Tuesday dummy	0.009 (0.009)	June dummy	0.046*** (0.011)	Chemicals dummy	0.705*** (0.014)	Public sector dummy	0.001 (0.034)
Wednesday dummy	-0.025*** 0.009	July dummy	0.002 (0.011)	Containers dummy	-0.149*** (0.013)	Software dummy	0.594*** (0.034)
Thursday dummy	-0.061*** 0.009	August dummy	-0.003 (0.011)	Electronics dummy	0.040*** (0.009)	Telecommun ications d.	0.205*** (0.028)
Friday dummy	-0.195*** (0.009)	September dummy	0.056*** (0.011)	Fashion dummy	0.324*** (0.011)	Transporta-ti on dummy	0.004 (0.018)
Saturday dummy	-0.968*** (0.008)	October dummy	0.041*** (0.011)	Food and beverage d.	0.494*** (0.013)	Travel dummy	0.699*** (0.027)
Sunday dummy	-1.053*** (0.008)	November dummy	0.034*** (0.012)	Healthcare dummy	0.226*** (0.015)	Constant	1.448*** (0.011)

Table 8.2. Full results obtained from the baseline firm-level regression with the daily number of transactions as the dependent variable

The dependent variable is the logarithm of the daily number of transactions for each seller-currency pair. Trade exposure is the percentage of trade with Russia, Belarus and Ukraine for a country, in proportion to its total trade. War dummy variable equals 1 for all dates on and after the first day of the war in Ukraine, 24-02-2022. Rolling sum is a past 30-day rolling sum of the number of transactions for a given seller-currency pair. Monday is the omitted baseline category of the day of the week dummy variable vector and January is the omitted baseline category of the month of the year vector. ‘Construction, Materials and Industrials’ is the omitted baseline category of the industry dummy variables. Names of the industry groups have been simplified for brevity.
Robust standard errors are reported in parenthesis. Coefficients significant at 10%, 5% and 1% are marked with *, ** and ***, respectively. Results are rounded to 3 decimal places.

8.3. Appendix C: Detailed results of individual country analysis

Buyers' country	War dummy coeffi-cient	War dummy st. error	Number of observa-tions	Buyers' country	War dummy coeffi-cient	War dummy st. error	Number of observa-tions
Australia	-0.066	0.052	11,409	Malaysia	-0.368***	0.086	762
Austria	0.104*	0.053	7,668	Mexico	0.425***	0.056	6,401
Belgium	0.169***	0.022	30,265	Morocco	0.139	0.138	192
Canada	0.168***	0.038	12,428	Netherlands	0.082***	0.021	52,291
Chile	-0.120*	0.071	512	New Zealand	0.548***	0.052	5,475
Costa Rica	-0.048	0.030	315	Norway	-0.012	0.042	4,955
Czechia	-0.239	0.189	293	Peru	0.324	0.231	450
Denmark	0.077**	0.036	8,674	Poland	0.209***	0.071	1,965
Dominican Republic	-0.073	0.141	733	Romania	-0.732***	0.272	381
Ecuador	-0.010	0.088	853	Singapore	-0.023	0.080	939
Finland	0.038	0.062	2,431	South Africa	0.361***	0.079	2,938
France	0.225***	0.052	3,561	Spain	0.245***	0.071	3,342
Germany	0.108***	0.030	22,804	Sweden	-0.057*	0.031	8,935
Guatemala	-0.012	0.141	210	Switzerland	0.194***	0.037	6,703
Hungary	-0.074	0.078	1,468	Thailand	-0.155	0.095	125
Ireland	0.056	0.039	7,566	Tunisia	-0.197	0.149	1,021
Israel	0.148*	0.085	762	United Kingdom	0.073***	0.022	41,585
Italy	0.508***	0.062	8,614	United States	0.151***	0.021	69,896
Latvia	-0.056	0.156	451	Uruguay	-0.239**	0.103	709

Table 8.3. Details of the coefficients of the dummy variable representing dates during the war obtained from regressions estimated separately for each country in the sample

The dependent variable is the logarithm of the daily number of transactions for each buyer-currency pair. War dummy variable equals 1 for all dates on and after the first day of the war in Ukraine, 24-02-2022. Rolling sum of transactions from the past 30 days, time and industry controls, as described in Chapter 5, are included in each regression. Regressions are estimated on data associated with one of the 38 buyer countries at the time.
Robust standard errors are reported in column 3. Coefficients significant at 10%, 5% and 1% are marked with *, ** and ***, respectively. Results are rounded to 3 decimal places.

9. **References**

Boehm, C. E., Flaaen, A., and Pandalai-Nayar, N. (2019). Input Linkages and the Transmission of Shocks: Firm-Level Evidence from the 2011 Tōhoku Earthquake. *The Review of Economics and Statistics, 101(1):* 60–75. DOI: 10.1162/rest_a_00750

Coibion, O., Gorodnichenko, Y., and Ropele, T. (2020). Inflation Expectations and Firm Decisions: New Causal Evidence. The Quarterly Journal of Economics, 135(1), 165–219. DOI:10.1093/qje/qjz029

Dunn, A., Hood K., Batch A., and Driessen A. (2021). Measuring Consumer Spending Using Card Transaction Data: Lessons from the COVID-19 Pandemic. *AEA Papers and Proceedings, 111:* 321-25. DOI: 10.1257/pandp.20211049

Ernst J. (2022, March 30). How Ukraine can win this war. NBC News, Politics & Policy. Retrieved from www.nbcnews.com/think/opinion/ukraine-can-win-war-russia-americas-lethal-aid-rcna22083

European Commission, European Statistical Office. (2022a). Sentiment indicators - monthly data. [Dataset] Retrieved from appsso.eurostat.ec.europa.eu/nui/show.do?dataset=ei_bssi_m_r2

European Commission, European Statistical Office. (2022b). *Glossary:Gross available energy.* Retrieved from ec.europa.eu/eurostat/statistics-explained/index.php?title=Glossary:Gross_available_energy

Faiola, A. (2022, March 30). How Ukraine could lose land but still win the war with Russia. The Washington Post, Today's Worldview. Retrieved from www.washingtonpost.com/world/2022/03/30/ukraine-partition-russia-deal-land/

Idrees Ali, I., and Stewart, P. (2022, April 06). Russia has completed withdrawal from around Kyiv -U.S. defense official. Reuters, World. Retrieved from www.reuters.com/world/us-assesses-russia-completes-withdrawal-around-kyiv-us-defense-official-2022-04-06/

Johnson, R. C. (2014). Trade in Intermediate Inputs and Business Cycle Comovement. *American Economic Journal: Macroeconomics, 6 (4):* 39-83. DOI: 10.1257/mac.6.4.39

Madhani, A., Knickmeyer, E., Merchant, N., and Jaffe A. (2021, December 04). US intelligence finds Russia planning Ukraine offensive. The Associated Press News. Retrieved from apnews.com/article/joe-biden-europe-russia-ukraine-vladimir-putin-ff5ac633e928a01b7d6d3bdbe2330fe2

McCann, F., and Myers, S. (2020). COVID-19 and the transmission of shocks through domestic supply chains. *Financial Stability Notes, Vol. 2020, No. 3. Central Bank of Ireland.* Retrieved from www.centralbank.ie/docs/default-source/publications/financial-stability-notes/no-3-covid-19-and-the-transmission-of-shocks-through-domestic-supply-chains-(mccann-and-myers)

Melitz, M. J. (2003). The Impact of Trade on Intra-Industry Reallocations and Aggregate Industry Productivity. *Econometrica, 71(6),* 1695–1725. Retrieved form www.jstor.org/stable/1555536

Data sources

Bank of International Settlements, BIS Statistics Warehouse. (2022). Exchange Rates. [Dataset] Retrieved from stats.bis.org/#ppq=XRU_M_AVG_BISSH_24M

European Commission, European Statistical Office. (2022c). Imports of natural gas by partner country. [Dataset] Retrieved from ec.europa.eu/eurostat/databrowser/view/NRG_TI_GAS__custom_2787083/default/table

European Commission, European Statistical Office. (2022d). Simplified energy balances. [Dataset] Retrieved from ec.europa.eu/eurostat/databrowser/view/NRG_BAL_S__custom_2787178/default/table

International Monetary Fund, International Financial Statistics. (2022). Prices, Production, and Labor selected indicators. [Dataset] Retrieved from data.imf.org/regular.aspx?key=63087884

World Bank, World Integrated Trade Solution. (2022). UN COMTRADE By Product. [Dataset] Retrieved from wits.worldbank.org/WITS/WITS/Default-A.aspx