The Real Impact of Sovereign Credit Ratings

Daniel Leff Yaffe

June 18, 2018

PRELIMINARY, DO NOT CITE

Abstract

This paper studies how CDS spreads react to changes in sovereign credit ratings. Since countries with increasing macroeconomic perspectives should show both a decline in their CDS spread and an increase in their credit rating, even if a credit change is not causal we will still observe correlation between these variables. To tackle this endogeneity problem I use the synthetic control method and identify treatment effects on 55 different sovereign credit ratings changes that occurred between 2008 and 2016. The results suggest that a credit upgrade decreases the CDS spread in the first 5 days by about 1 basis points. In contrast, a credit downgrade increases the CDS spread in the first 5 days by about 5 basis points. For both upgrades and downgrades, the possibility that this effect is not temporary can't be rejected.

1 Introduction

When two parties enter a credit default swap (CDS) contract the buyer makes a periodic payment, called the CDS spread, and in return receives a payoff from the seller if some underlying loan goes into default. As such, CDS spreads are indicators of the market perceived risk of an asset: if loan A has a lower CDS spread than loan B, then we infer that loan A has a smaller perceived probability of going into default.

Another indicator of the ability bond issuers have of paying back a loan are credit ratings. These are ordinal rankings published by private companies called Credit Rating Agencies (CRAs). The three most important CRAs, which control about 95% of the credit ratings market (Kingsley, 2012), are Standard & Poors (S&P), Moody's, and Fitch. Across these agencies the rating system is a little bit different, however, as Table 1 illustrates, they are easily comparable. It is important to note that these rankings have a certain degree of subjectivity; so even though they are usually very similar, they need not be equal.

This paper focuses on sovereign debt, and studies whether changes in sovereign credit ratings (SCR) can affect the CDS spreads of their respective countries. That is, can a credit change announcement from a CRA affect the market perceived risk of a sovereign bond? If SCR have an impact on the CDS spreads then the financing of such debt would be affected by this external rating. In such case countries would be interested in trying to influence CRAs since 1 basis point can represent millions of dollars. One example of such behavior was observed in 2011, when U.S. Treasury Secretary Timothy Geithner tried to persuade S&P not to downgrade the rating on U.S. bonds (Gill & Gill, 2015). His attempts were ultimately unsuccessful and S&P downgraded the U.S. to AA+ on Aug. 05, 2011.

The difficulty in assessing whether SCR have a casual impact on CDS spreads lies in the endogeneity between both variables. Since countries with increasing macroeconomic perspectives are more able to pay off their debts, they should simultaneously show a decline in their CDS spread and an increase in their credit rating. Therefore, even if credit changes have no causal impact on CDS spreads we will still observe a correlation between these variables.

Table 1: Comparison of Ratings across the
Big Three Credit Rating Agencies

Moody's	S&P	Fitch	Rating description
Aaa	AAA	AAA	Prime
Aa1	AA+	AA+	High grade
Aa2	AA	AA	
Aa3	AA-	AA-	
A1	A+	A+	Upper medium grade
A2	A	A	
A3	A-	A-	
Baa1	BBB+	BBB+	Lower medium grade
Baa2	BBB	BBB	
Baa3	BBB-	BBB-	
Ba1	BB+	BB+	Non-investment grade speculative
Ba2	BB	BB	
Ba3	BB-	BB-	
B1	B+	B+	Highly speculative
B2	B	B	
B3	B-	B-	
Caa1	CCC+	CCC+	Substantial risks
Caa2	CCC	CCC	
Caa3	CCC-	CCC-	
Ca	CC	CC	Extremely speculative
	C	C	Default imminent
C	RD	DDD	In default
/	SD	DD	
/	D	D	

To see this correlation, Figure 1 shows a scatter plot with cross-section data on 66 countries for September 19, 2016. The plot relates the 5-year CDS spread on sovereign debt and S&P's long-term rating for sovereign debt issued on foreign currency. As expected, a negative relationship is found: countries with higher ratings tend to have lower CDS spreads meaning their probability of defaulting is lower. For this sample Spearman's rank correlation coefficient is equal to -0.93.

Figure 1: Scatter plot for 66 countries, Sept. 19, 2016

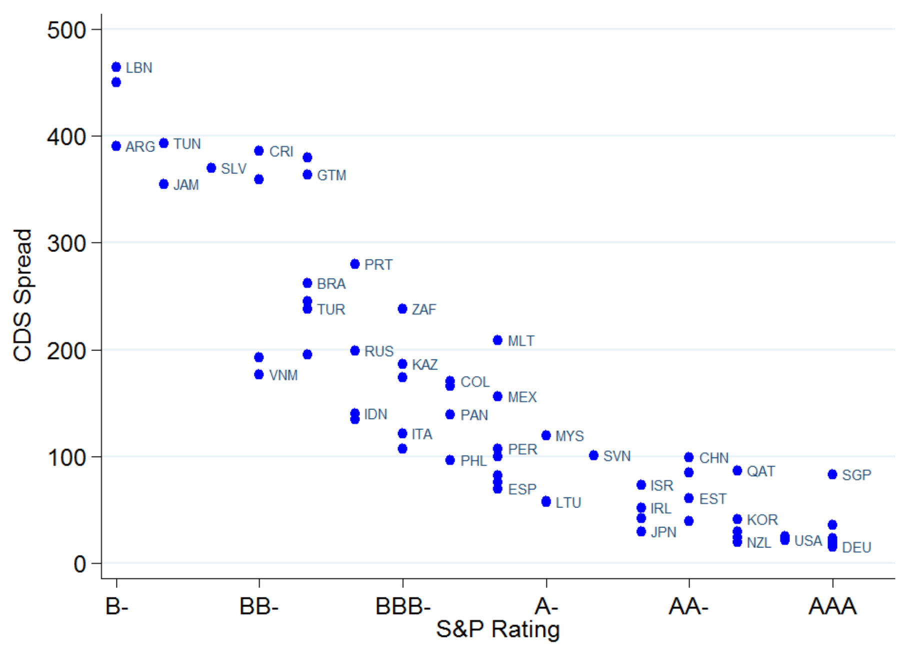

Note: CDS spreads data is from Thomson Reuters Datastream. Ratings are for long-term debt in foreign currency.

So under what circumstances can we expect credit ratings to have a casual effect on CDS spreads? Clearly if the public had perfect information, and there were no frictions in the economy, we would not expect any casual effect. In such a world one could think of CRA as taking publicly available information and simply transforming it to a different scale. However, if CRAs have private information about sovereigns, or even if they don't but economic agents believe they do, then a change in a SCR will signal new information which might end up having an impact. In reality, even though a country's problems might already be perceivable to the world, a credit downgrade stands out as additional bad news. So it is not completely unreasonable for the downgrade to affect the country even more.

Another channel through which SCR can have an effect in the economy is through institutional frictions; for example regulation or internal rules might force some investors to only hold bonds with a certain minimum credit rating. Perhaps the most important institutional friction is a consequence the standardized approach; a method set forth by the Basel II capital adequacy rules for banking institutions. Under this approach banks are required to use ratings from CRAs to risk-weight their assets when computing their capital adequacy ratio. Table 2 shows the risk weight assigned on sovereign debt for different levels of credit ratings. The ratings are expressed in the S&P's scale for notation purposes only; banks are free to chose any certified CRA whose ratings they want to use in risk-weighting their assets.[1]

Table 2: Risk-Weights on Sovereign Bonds under Basel's II Standardized approach

Credit Assessment	**Risk Weight**
AAA to AA-	0%
A+ to A-	20%
BBB+ to BBB-	50%
BB+ to B-	100%
Below B-	150%
Unrated	100%

[1]Importantly, banks under the standardized approach must disclose which CRAs they are using, and if they decide to work with more than one they must use the highest of the risk weights for each asset they hold. This provision ensures that banks are not able to choose multiple CRAs in order to achieve lower risk weights (Roy, 2005).

In order to study the impact of a SCR change on its respective CDS spread, I construct counterfactuals using the Synthetic Control (SC) method (Abadie et al., 2010). Section 2 of the paper goes over the data used and Section 3 explains the methodology. In Section 4 I present and explain the results. Finally Section 5 concludes.

2 Data

I use daily end-of-day data[2] on CDS spreads from Thomson Reuters Datastream. Data from 69 countries is obtained (a list of these is included in Appendix A. For some countries the time period covers from Dec. 14, 2007 until May 1, 2017. For other countries the time period is shorter, so I end up we an unbalanced panel data.

The other piece of data used are SCR and outlooks on countries' ratings. I use ratings from S&P because research has shown that these ratings usually precede those from other CRAs, that they are more frequent and less anticipated (Gande & Parsley, 2005; Reisen & von Maltzan, 1999). S&P publishes historical data on every credit announcements issued for sovereign ratings on their Global Credit Portal website. Using this information I construct a database, with daily frequency, containing the SCR and credit outlook for each of the 69 countries for which I have CDS data.

Columns 2 through 4 of Table 3 show, per initial rating level, the number of credit upgrades and downgrades presented withing the sample.

[2] Any mention of days should be understood as business days.

Table 3: Number of Credit Changes in our Sample

Initial Rating	Total in sample			Remaining after filters		
	Upgrades	Downgrades	Both	Upgrades	Downgrades	Both
CCC-	0	1	1	0	0	0
CCC	0	2	2	0	0	0
CCC+	5	5	10	0	0	0
B-	6	6	12	0	0	0
B	3	8	11	0	0	0
B+	3	6	9	0	0	0
BB-	6	3	9	0	0	0
BB	9	6	15	6	0	6
BB+	10	7	17	5	2	7
BBB-	10	7	17	8	1	9
BBB	6	13	19	6	7	13
BBB+	4	10	14	1	3	4
A-	4	6	10	1	0	1
A	8	9	17	3	2	5
A+	3	5	8	2	1	3
AA-	1	3	4	1	0	1
AA	1	4	5	0	3	3
AA+	2	5	7	0	2	2
AAA	0	7	7	0	1	1
Total	81	113	194	33	22	55

3 Methodology

To study how changes in SCR affect CDS spreads I use the SC method (Abadie et al., 2010). Let i index countries and $\{x_i\}_{i=0}^N$ be the list of countries. First, I explain how this method is used to identify the impact of a credit change country x_0 experiences at date $t = \tau$. Thus, I define x_0 as the country receiving treatment and $\{x_i\}_{i=1}^N$ as the donor pool. Treatment is defined as receiving a credit change, whether it is an upgrade or a downgrade. However, in section 4 I allow for asymmetric impacts of upgrades and downgrades.

Let c_{0t} be the CDS spread of the treatment unit and $\{c_{it}\}_{i=1}^N$ be the CDS spreads of every other country at day t. The SC method will deliver a set of weights $\{w_i\}_{i=1}^N$ such that $\sum_{i=1}^N w_i = 1$, $w_i > 0 \ \forall i$ and the linear combination

$$\hat{c}_{0t} = \sum_{i=1}^{N} w_i c_{it} \text{ for } t < \tau$$

mimics the the behavior of the treatment group.

$\hat{c}_{0t}$ is then called the SC. The optimization procedure, which minimizes the root mean squared error (RMSE) for a selected number of periods prior to the intervention usually delivers only positive weights for a small number of countries. Importantly, these weights are obtained in an objective manner, so a researcher can't influence the result to go in a specific direction.

To identify the impact of a credit change the identifying assumption is that in the absence of treatment the weights $\{w_i\}_{i=1}^N$ would remain constant. As such, these weights may be used to predict what would have happened to the CDS spread of country x_0 had it not been subject to the credit change. This counterfactual, which is simply the SC for dates on and after the treatment, can be easily computed as

$$\hat{c}_{0t} = \sum_{i=1}^{N} w_i c_{it} \text{ for } t \geq \tau$$

Let j index credit changes. Then using the SC method the impact Δ_{ijt} of credit change j, for a country indexed with i, at date t, is obtained by computing

$$\Delta_{ijt} = c_{ijt} - \hat{c}_{ijt}$$

So far I have explained how to apply the SC method to a single credit change. In this paper I apply the SC method individually to 55 selected credit changes observed withing the sample.[3] Then, I put together all our estimated impacts Δ_{ijt} and obtain an empirical distribution of treatment effects Λ_h for $H+1$ horizons:

$$\{\Lambda_h\}_{h=0}^{H} \equiv \left\{ \left\{ \{\Delta_{ijt}\}_{i \in I} \right\}_{j \in J} \right\}_{t=\tau_{ij}}^{\tau_{ij}+H}$$

That is, I obtain an empirical distribution of impacts to credit changes from τ (effect on impact) until $\tau + H$ (effect H business days after impact). In section 4, I fix H to be equal to 20 days and study the median and the mean of such distributions. These statistics are useful since they are more representative indicators of the impact of credit changes than a single credit change alone.

3.1 Techinal Considerations

In order to apply the SC control method one has to define the treatment group, the donor pool and over what pre-treatment period the RMSE should be minimized. This subsection goes over some technical considerations when applying the method to CDS data. I explain criteria for minimizing the RMSE, situations under which certain credit changes are excluded from the analysis, restrictions to the donor pool, and other important considerations.

After applying all of the filters outlined below the total number of credit changes that can be studied drops from 194 to 55, as shown in the last 3 columns of Table 3. Of these

[3] The criteria for to select a credit change is explained in subsection 3.1.

credit changes the first one is dated on Nov. 27, 2008, while the last one on Feb. 17, 2016.

For these subsections let country x_0 receive a credit change at date τ. Moreover, we r_0^L and r_0^H denote the low and the high rating that country x_0 experienced when it was subjected to the credit change.

3.1.1 Minimizing the RMSE and in-sample fit

I use CDS data from the donor pool in each of the 10 business days prior to the credit change to minimize the RMSE of the pre-treatment period. Also, since the SC method does not adjust for an intercept, I apply the method to CDS data demeaned by the average CDS spread of the 10 days prior to treatment. This adjustment is specially beneficial for situations when there are CDS spreads in the donor pool that mimic the treatment unit well, but all of them lie above (or below) the treatment unit.

It is important to make sure that the SC behaves similarly to the treatment unit during the pre-treatment period. With this in mind I only consider those SC that deliver a RMSE lower than 5. The intuition behind this criteria is that we do not want to make out-of-sample forecasts if our in-sample fit is poor.

3.1.2 Valid Treatment Units

To avoid contamination of multiple credit changes I do not consider a credit change of country x_0 that occurred at date τ if the same country received another credit change 30 days before or after τ. I also exclude a credit change if there was an outlook change 30 days before or after τ.

I also only consider "unit" credit changes, i.e. the rating change must have been of a single step in the ranking list. For example, Cyprus' credit change of Jan. 13, 2012 is not considered because it downgraded Cyprus' SCR from BBB to BB+, without passing through BBB-.

Finally, only changes where $r_0^L \succcurlyeq$BB are analyzed. That is, I don't consider upgrades that started with BB- nor downgrades starting with BB. On the margin, this drops 7 upgrades and 2 downgrades from our sample. This condition is applied because CDS spreads on countries with very low ratings are too erratic to be useful.

3.1.3 Restrictions to the Donor Pool

Only countries that did not receive a rating or outlook change in the 30 days before and after τ enter the donor pool. Additionally, I exclude countries whose rating was too far away from that of the treatment unit. Specifically, I only consider countries that maintained a rating between 3 units above r_0^H and 3 units below r_0^L. For example if country x_0 had a downgrade from A to A+, then I only allow countries with ratings between BBB and AA+ into the donor pool. This restriction does not affect results since the SC method usually finds that the countries that best predict the CDS spread of x_0 are those with similar ratings. In fact, the main advantage of including this condition is improving the speed with which the SC optimization is performed.

3.1.4 Other considerations

Poor Quality Data: For a few countries, the CDS spread data is of very poor quality. The data quality may be observed by noting that the CDS spread remains constant for many days, without even moving a decimal point. This is probably a consequence of there being few CDS trades for said countries, so instead of observing the true market price of CDS contracts the same price as the day before is reported. Therefore, I apply the following criteria by which I drop low quality CDS data from the sample: if the CDS of country x_i had a sample standard deviation over the last 10 business days of less than 0.50 then I drop the CDS spread of said country, at such date, from the data set. This criteria drops 21% of the CDS spreads data.

Missing Observations on CDS Spreads: A country can only be the treatment unit or belong into the donor pool if it does not have missing CDS spread information or poor quality data (as defined above) for at least 10 days before and 20 days after treatment occurred. The 10 days before treatment condition is necessary to run the SC method, and the 20 days after is needed in order to make comparisons between the observed outcome and the counterfactual.

4 Results

First, I present a motivating example, in which I apply the CDS methodology to analyze Peru's credit upgrade on August 30, 2011. Then, I present results drawn from studying the distributions of impacts across the selected 55 credit changes. I study impacts of upgrades and downgrades separately, since asymmetric impacts between the two type of events have been suggested.

4.1 Motivating Example

To motivate using the synthetic control method, consider the SCR upgrade from BBB- to BBB that Peru received on August 30, 2011. The black line in Figure 2 shows how the CDS Spread of Peru behaved the 10 business days prior to the upgrade. For this case the donor pool consists of 26 countries. From these, the SC method gives only positive weights to 5 countries: Colombia (57.8%), Mexico (21.6%), Croatia (11.8%), Iceland (8%) and Portugal (1%). Intuitively, Colombia and Mexico are countries someone doing a Difference-in-Difference analysis would consider as potential controls even before looking at the data. The advantage of using SC is that there is no need to think about what countries would be good controls for each of the credit changes. Instead, the method automatically finds countries with similar pre-treatment behavior in their CDS spreads.

Figure 2: Before Peru's Credit Upgrade
of Aug. 30, 2011

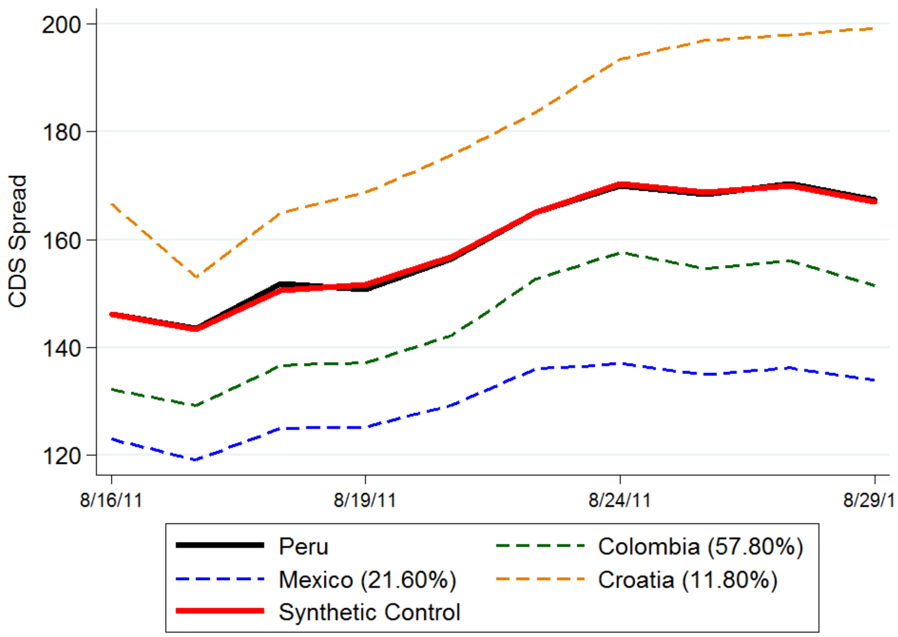

Figure 2 also shows us the CDS spreads for Colombia, Mexico and Croatia (which together account for 91.2% of the SC weights),[4] as well as the SC itself. As can be seen the in-sample predictive power is good in this example: Peru's SC basically overlaps with Peru's actual CDS spread.

In Figure 3 I use the same set of weights to create the counterfactual for the date the credit change occurred as well as the 20 business days that follow it. This example suggests that starting 5 days after the credit change Peru's CDS spread was about 5 basis points lower than what it would have been if S&P had not increased its rating

[4]These 3 series are scaled up and down for illustration purposes, i.e., their mean was adjusted for an easier visualization.

Figure 3: Before and After Peru's Credit Upgrade
of Aug. 30, 2011

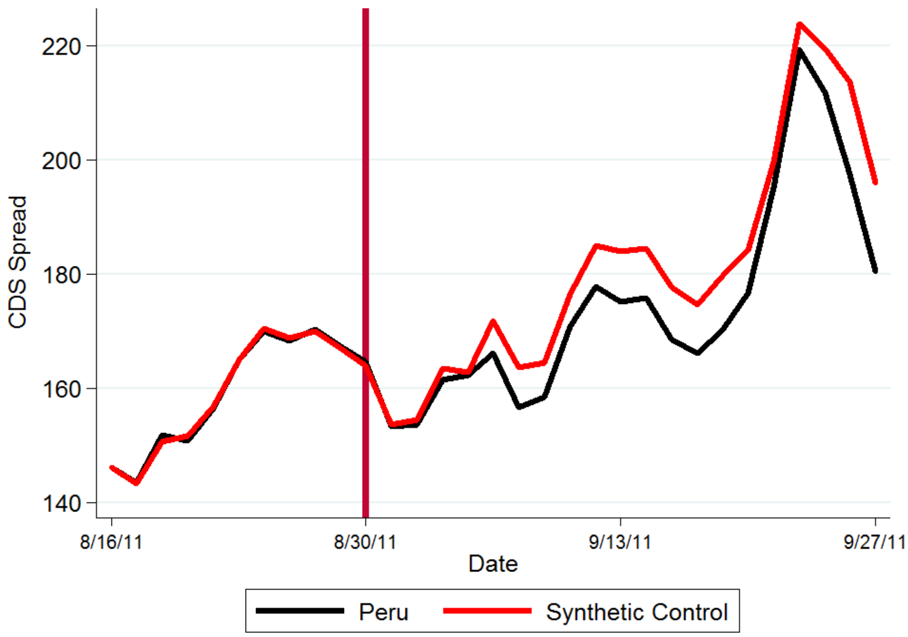

4.2 Empirical Distributions

The motivating example measures the impact of a single credit upgrade. Importantly, we do not know how representative this upgrade is of other credit changes. This is why in this subsection I create an empirical distributions $\{\Lambda_h\}_{h=0}^{20}$ that puts together results for a total of 33 upgrades and 22 downgrades. This allows computing medians and means that are more representative of the impact of credit changes than a single credit change alone.

Furthermore, by performing hypothesis tests on the median of the distribution, the question of whether credit changes have an impact on CDS spreads can be answered. The results suggest that they do.

Figures 4 and 5 show estimated impacts upon a credit upgrade and a credit downgrade, respectively. The time period is standardized so impacts occur at day 0. Days -10 to -1 constitute the the pre-treatment period used to minimize the RMSE in the SC method. I refer to dates 0 to 20 as the post-treatment dates.

Each figure shows the mean, median and interquartile range (IQR) of the estimated empirical distribution for the different horizons considered. From the 10 days before impact we can observe that, for the credit changes considered, the in-sample predictive power of the SC method is good. The mean and median are very close to zero, and the IQR is very narrow.

Looking at the medians and means of the post-treatment period, I find that the estimated impacts of upgrading and downgrading ratings have the expected signs. A credit upgrade decreases the CDS spread in the first 5 days by about 1 basis points. In contrast, a credit downgrade increases the CDS spread in the first 5 days by about 5 basis points. The results suggest that credit downgrades have stronger effects. The mean and median in Figure 5 are persistently farther away from zero than those in Figure 4. Also, by looking at the IQR, one may notice that for downgrades the distribution at each post-treatment horizon has a lot of more positive entries relative to the negative entries that follow a credit upgrade.

Needless to say, finding the adequate sign is not enough. It is also important to assess if these impacts are significant. To do so, I apply sign tests to see whether the medians of the estimated distributions are significantly different than zero. Table 4 shows the sign test for horizons 0 to 20. For upgrades the interest lies in whether the median impact is negative, so the adequate alternative hypothesis is H_A : Median of $\Gamma_t < 0$. In the other hand, for downgrades the alternative hypothesis is H_A : Median of $\Gamma_t > 0$.

The results show significant deviations from zero in both upgrades and downgrade. Moreover, the possibility that this effect is not temporary can't be rejected. Using a significance level of 10%, I reject the null of the median being equal to zero at horizons 2, 4, 19 and 20 for upgrades. Using the same significance level for downgrades, the test rejects the null at horizons 4, 5 9, 10, 11, 12, 13, 14, 16 and 17. These results also confirm a stronger effect from credit downgrades.

Figure 4: Effect of a Credit Upgrade

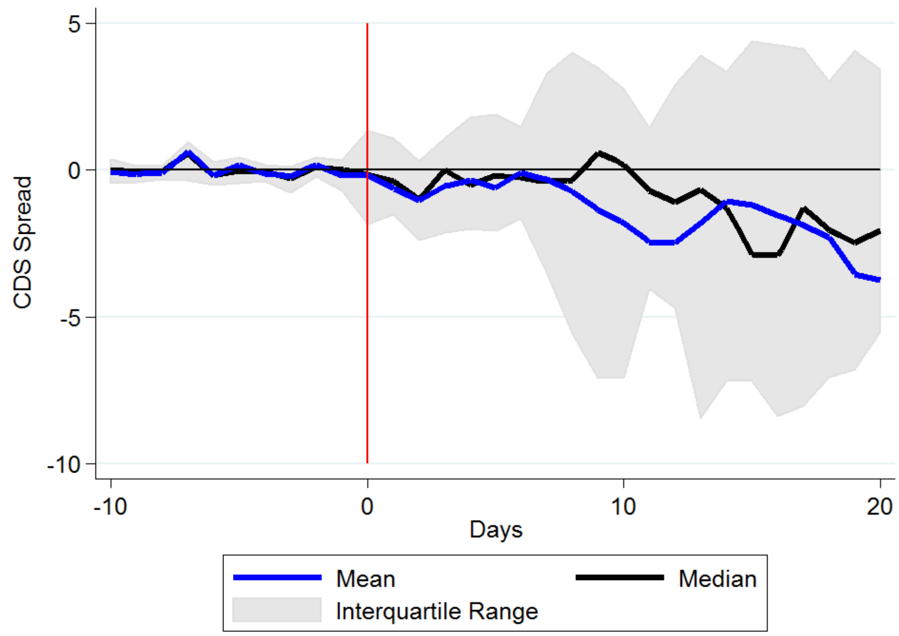

Figure 5: Effect of a Credit Downgrade

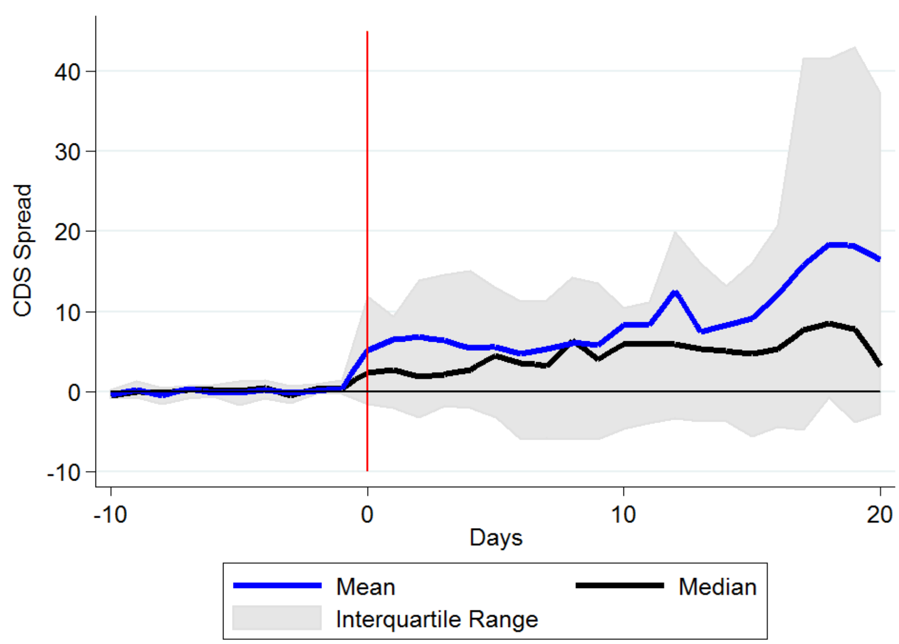

Table 4: Post-Treatment P-Values from Hypothesis Testing on the Median

Days After Credit Change	**Upgrades** H_A: Median of $\Gamma_t < 0$	**Downgrade** H_A: Median of $\Gamma_t > 0$
0	0.24	0.14
1	0.36	0.26
2	0.04	0.42
3	0.64	0.14
4	0.08	0.07
5	0.36	0.03
6	0.36	0.42
7	0.36	0.14
8	0.36	0.14
9	0.76	0.07
10	0.64	0.07
11	0.24	0.07
12	0.15	0.03
13	0.24	0.03
14	0.15	0.07
15	0.08	0.14
16	0.15	0.02
17	0.36	0.02
18	0.15	0.15
19	0.08	0.15
20	0.04	0.15

For robustness checks, Table 5 also checks whether the median in the pre-treatment period is different than zero. Here the adequate alternative hypothesis is H_A : Median of $\Gamma_t \neq 0$ for both upgrades and downgrades. For the pre-treatment period the null hypothesis can't be reject for any horizon (neither for upgrades nor downgrades). This gives confidence on the results from Table 4 as the pre-treatment period is not showing any predisposition towards being more positive or negative.

Table 5: Pre-Treatment P-Values from Hypothesis Testing on the Median

Days Before Credit Change	**Upgrades** H_A: Median of $\Gamma_t \neq 0$	**Downgrade** H_A: Median of $\Gamma_t \neq 0$
1	1.00	0.13
2	0.49	0.13
3	0.16	0.29
4	0.73	0.13
5	1.00	1.00
6	0.49	0.83
7	0.30	0.52
8	0.73	0.83
9	0.73	1.00
10	0.49	0.52

5 Conclusions

To study the impact of a SCR change on its respective CDS spread I construct counterfactual CDS spreads using the SC method. Each counterfactual allows us to analyze what would have been the CDS spread of a country receiving a credit change had it not been subject to it. Then, I construct and analyze empirical distributions of the impact of upgrading and downgrading such ratings at different horizons. The results suggest that a credit upgrade decreases the CDS spread in the first 5 days by about 1 basis points. In contrast, a credit downgrade increases the CDS spread in the first 5 days by about 5 basis points. For both upgrades and downgrades the possibility that this effect is not temporary can't be rejected.

Unlike previous work in this area, I make a serious attempt to deal with the endogeneity between SCR and CDS spreads and find interesting and significant results. I show there is a much stronger impact of credit downgrades and find reasonable estimates. Furthermore, robustness checks for the pre-treatment period give more confidence on the validity of the results; the SCs show no predisposition towards taking more positive or negative values before the credit change happens.

Bibliography

Abadie, A., Diamond, A., & Hainmueller, J. (2010). Synthetic control methods for comparative case studies: Estimating the effect of California's tobacco control program. *Journal of the American statistical Association*, 105(490), 493-505.

Gande, A., & Parsley, D. C. (2005). News spillovers in the sovereign debt market. *Journal of Financial Economics*, 75(3), 691-734.

Kingsley, P. (2012, Feb. 15). How credit ratings agencies rule the world. *The Guardian.*

Reisen, H., & von Maltzan, J. (1999). Boom and bust and sovereign ratings. *International Finance*, 2(2), 273-293.

Roy, P. V. (2004). Credit ratings and the standardised approach to credit risk in Basel II, Finance 0509014. *Economics Working Paper Archive EconWPA*, 1561-0810.

A Appendix: List of Countries in the Sample

Argentina	Hong Kong	Peru
Australia	Hungary	Philippines
Austria	Iceland	Poland
Bahrain	Indonesia	Portugal
Belgium	Iraq	Qatar
Brazil	Ireland	Romania
Bulgaria	Israel	Russia
Chile	Italy	Serbia
China	Jamaica	Singapore
Colombia	Japan	Slovakia
Costa Rica	Kazakhstan	Slovenia
Croatia	Korea	South Africa
Cyprus	Latvia	Spain
Czech Republic	Lebanon	Sweden
Denmark	Lithuania	Thailand
Dominican Republic	Malaysia	Tunisia
El Salvador	Malta	Turkey
Estonia	Mexico	Ukraine
Finland	Morocco	United Kingdom
France	Netherlands	United States
Germany	New Zealand	Uruguay
Greece	Norway	Venezuela
Guatemala	Panama	Vietnam