WP/16/117

IMF Working Paper

The Impact of Trade Agreements: New Approach, New Insights

Swarnali Ahmed Hannan

INTERNATIONAL MONETARY FUND

 WP/16/117

IMF Working Paper

Strategy, Policy and Review Department

The Impact of Trade Agreements: New Approach, New Insights

Prepared by Swarnali Ahmed Hannan

Authorized for distribution by Martin Kaufman

June 2016

Abstract

The Trans-Pacific Partnership (TPP) has reinvigorated research on the ex-ante impact of trade agreements. The results from these ex-ante models are subject to considerable uncertainties, and needs to be complimented by ex-post studies. The paper fills this gap in recent literature by employing ***synthetic control methods (SCM)*** – currently extremely popular in micro and macro studies – to understand the impact of trade agreements in the period 1983–1995 for 104 country pairs. The key advantage of using SCM to address selection bias – one of the persisting issues in trade literature – is that it allows the effect of unobserved confounder to vary with time, as opposed to traditional econometric methods that can deal with time-invariant unobserved country characteristics. Using SCM approach, the paper finds that trade agreements can generate substantial gains, on average an increase of exports by 80 percentage points over ten years. The export gains are higher when emerging markets have trade agreements with advanced markets. The paper shows that all the countries in NAFTA have substantially gained due to NAFTA. Finally, there is some evidence that trade agreements can potentially lead to slight import diversion, but not export diversion.

JEL Classification Numbers: F1, F13, F14

Keywords: Trade agreements, international trade flows, synthetic control method.

Author’s E-Mail Address: sahmed@imf.org

Contents

I. INTRODUCTION[1]

The Trans-Pacific Partnership (TPP), the largest preferential trade agreement to date, has reinvigorated interest on understanding the impact of trade agreements. Recent academic research has thus far focused on ex-ante analysis of trade policy changes, using computable general equilibrium model (CGE), to understand the impact of TPP. Due to the considerable uncertainties involved in ex-ante models deciphering the future impact of forthcoming trade agreements, there needs to be a renewed effort in understanding the ex-post impact of past trade agreements. This is also crucial since ex-ante models have typically under-estimated the impact of trade agreements with the actual outcome turning out to be significantly better than model predictions (Kehoe, 2003).

This paper fills the gap in recent literature by employing synthetic control methods (SCM) – currently very popular in micro and macro studies – to understand the impact of past trade agreements. Synthetic control method (SCM) is a recent econometric tool for comparative studies where the choice of comparison unit is a systematic data driven procedure. SCMs create synthetic (artificial) control or comparison units based on their similarity to the treated country before the treatment actually takes place. This is done based on a weighted average of past observable covariates and past realizations of the outcome variables. The evolution of the actual outcome of the treated unit post- treatment is then compared against the outcome of the synthetic unit, and the difference is interpreted as the treatment effect.

Why employ SCM on trade agreements? The gravity equation is typically used to explain cross-sectional variation in country pairs' trade flows in terms of the countries' incomes, bilateral distance, and dummy variables for common languages, common land borders, etc. To understand the impact of trade agreements, literature usually uses a dummy variable in the right hand side for the presence of trade agreements. Earlier studies, assuming that the trade agreement was an exogenous variable, found no meaningful impact. However, later studies showed that trade agreements, in practice, are not exogenous random variables,

[1] I am thankful to Tamim Bayoumi, Diego Cerdeiro, Valerie Cerra, Varapat Chensavasdijai, Saiful Hannan, Christian Henn, Martin Kaufman, Brad McDonald, Michele Ruta, and seminar participants at IMF Trade Seminar for their helpful comments and suggestions. All remaining errors are mine.

but countries are likely to select endogenously into trade agreements for reasons not observable but correlated with the level of trade, hence biasing the results. Literature subsequently used various econometric techniques to address endogeneity problem arising due to selection bias. The key advantage of SCM in addressing this endogeneity problem associated with trade agreements is that it allows the effect of unobserved confounders to vary with time, as opposed to traditional econometric methods that can deal time-invariant unobserved country characteristics only. This feature has also made SCM popular in other fields. Estimating trade impacts using SCM would thus be a novel way of addressing the endogeneity bias problem in trade literature.

Using SCM approach to the gravity model, this paper looks at 104 country pairs that had trade agreements in the period 1983–1995. The paper finds that trade agreements can generate substantial gains, on average an export boost of 80 percentage points over ten years. The annual export growth is 3.8 percentage points higher due to trade agreements. Overall, the export gains are even higher when the anticipation effect – trade increasing before the agreement is enacted – is accounted for. While previous studies have not differentiated trade gains across income groups, this paper shows that trade gains can be substantially higher for emerging markets when they have a trade agreement with advanced markets. The gains are relatively less, compared to the average, for advanced economies exporting to emerging markets. The results also show that, contrary to some of previous key studies, each bilateral country pair involving Canada, Mexico and the U.S. has generated more exports due to NAFTA. The average exports have increased by 79 percentage points over ten years due to NAFTA, with annual export growth 4.3 percentage points higher. Finally, the paper touches upon the issue of trade diversion – is the export gains at the expense of countries that are not part of trade agreement? The paper exploits the underlying logic of SCM to look at trade diversion in a novel manner. For each country engaged in a trade agreement, SCM is employed to the pair representing the country's trade with its top trading partners (top exporter and top importer separately) that are outside the trade agreement. The synthetic unit in this scenario represents the counterfactual of how trade would have evolved with the third (non-signatory) country under the absence of trade agreement. The results show evidence of slight import diversion, but not export diversion.

This paper contributes to the strand of ex-post trade literature that attempts to address the endogeneity issue due to the selection bias when looking at the impact of trade agreements (see Section II on literature review). As far as the author is aware, this is the first attempt to employ SCM across a large number of trade agreements. In addition, this study takes a more granular look into how trade agreements affect different income groups. The paper also exploits the SCM methodology in a novel manner to gather insights about trade diversion.

The rest of the paper is organized as follows. Section II reviews literature, both on ex-ante and ex-post assessments, to get a sense of context, and then provides motivation for this paper. Section III discusses the theoretical details underlying SCM methodology. Section IV discusses the results on trade creation and trade diversion, including the placebo tests. Section V concludes.

II. Literature Review and Motivation

Academic literature on the impact of trade agreements can be bucketed into two groups: i) *ex-post assessment*, employing standard regressions, on historical data; and ii) *ex-ante predictions,* using computable general equilibrium (CGE) methods, based on projections about future.

Ex-post studies:

Over the past 40 years, the 'gravity equation' has been used as the empirical workhorse in international trade to study the *ex-post* effects of trade agreements. As described by Baier and Bergstrand (2007), the gravity equation is typically used to explain cross-sectional variation in country pairs' trade flows in terms of the countries' incomes, bilateral distance, and dummy variables for common languages, common land borders, and for the presence of trade agreements. UNCTAD/WTO (2012) has detailed exposition of the main underpinnings of this model. The earlier ex-post studies based on gravity equation for international trade flows show no clear and convincing evidence about the impact of trade agreements in boosting trade (Tinbergen, 1962; Bergstrand, 1985). Furthermore, Ghosh and Yamarik (2004) use extreme-bounds analysis to test the robustness of trade agreement dummy coefficients and find that the coefficients are ‘fragile’, indicating that the results are not very robust.

One of the key reasons responsible for earlier studies showing little impact is that trade agreements are represented in the right-hand-side as exogenous dummy variables. As highlighted in more recent work, trade agreements, in practice, are not exogenous random variables, but countries are likely to select endogenously into trade agreements for reasons not observable but correlated with the level of trade, hence biasing the results. Magee (2003) shows that countries are more likely to be preferential trading partners if they have significant bilateral trade, are similar in size, and are both democracies. Baier and Bergstrand (2004) find strong cross-section empirical evidence that pairs of countries that have free trade agreements (FTA) tend to share economic characteristics that their theory suggests should enhance the net economic welfare gains from an FTA for the pairs’ representative consumers.

Academic literature has subsequently proceeded to employ econometric techniques to address endogeneity problem arising due to selection bias. Baier and Bergstrand (2007) and Magee (2003) use instrumental variables with cross section data to address this bias. Baier and Bergstrand (2009) used nonparametric (matching) econometric techniques, while other studies have used Probit models (Mansfield and Reinhardt, 2003 and 2008; Baier and Bergstrand, 2004). These studies show larger impact when controlling for endogeneity, however results are sensitive to the years considered, variables included and estimation techniques. For example, Baier and Bergstrand (2009, 2007) find that trade agreements increase trade by 100 percent after ten years.

When estimating the impact of trade agreements, a related question is where trade agreements lead to diversion in trade, the idea that members of trade agreements will redirect their trade amongst themselves to take advantage of preferential treatment. This could be at the expense of other countries that are not part of trade agreements and potentially cause misallocations if non-member products were more competitive in absence of preferences. Early theories of free trade agreements emphasized trade diversion effects (Viner 1950, Lipsey 1960). Empirical studies since then recognize that economies with significant pre-agreement trade are “natural trading blocs” and their agreements are likely to lead to more trade creation than trade diversion (Frankel, Stein, and Wei 1995). More recently, however, studies have shown evidence of trade diversion. Carrere (2006) finds, studying seven regional trade agreements, that most of the trade agreements resulted in an increase in intra-regional trade beyond levels predicted by the gravity model, often coupled with a reduction in imports from the rest of the world, and at times coupled with a reduction in exports to the rest of the world, suggesting evidence of trade diversion.

Ex-ante studies:

While the gravity model can be used to study the impact of trade agreements that has been in place for a sufficient period of time for its effects to be observable in the data, partial equilibrium (PE) or general-equilibrium (GE) models are usually used for the ex-ante analysis of trade policy changes. As summarized in UNCTAD/WTO (2012), the idea is to numerically simulate the general equilibrium structure of the economy. Two scenarios are constructed: i) a baseline scenario with no policy change, ii) another scenario which models the policy changes implicit in the trade agreement. The difference between the two scenarios is then the impact of the trade agreement. Recently, the Trans-Pacific Partnership (TPP) has reinvigorated research using this strand of literature (Petri and Plummer, 2016; Aichele and Felbermayr, 2016; Ciuriak and Xiao, 2014).

While this approach is extremely useful in gauging the impact of trade agreements before they materialize, the results are surrounded by considerable uncertainty (Petri and Plummer, 2016) due to the difficulty of quantifying all the channels responsible for boosting trade due to trade agreements. CGE models rely on a complex network of assumptions, and the results may change substantially with a small change in the assumed framework (Hufbauer and Schott, 2005). Making some assumptions is unavoidable in translating the agreement into a numerical framework to encompass all the drivers which can lead to considerable uncertainty, as pointed out by Petri and Plummer (2016):

> ” The most important data points of the model include trade and investment barriers for each product on each exporter-importer link. These are difficult to estimate because some impediments are hard to pinpoint and because complex patterns of existing bilateral trade agreements affect much intra-TPP trade. Information on tariffs is reasonably complete and reliable, but data on NTBs, which are more significant, are measured less accurately and leave gaps to be filled.”
>
> Petri and Plummer (2016)

Compared to ex-post studies, ex-ante studies have generally found lower impact of trade agreements. Kehoe (2003), comparing predictions of three ex-ante CGE studies of NAFTA with observed outcomes, find that the trade increases in most sectors outstripped the predictions more than ten-fold. Corcos, del Gatto, Mion and Ottaviano (2012) find similar results for the EU.

Motivation:

The motivation for this paper comes from the renewed interest on the impact of trade agreements due to TPP. The TPP, the largest preferential trade agreement to date, is a major development in the trade landscape, and has generated immense interest. Consequently, in recent years, trade literature has focused on understanding the impact of TPP employing CGE models. However, due to the uncertainty of the results from CGE models, this needs to be complimented by ex-post studies of what happened during past trade agreements. This paper contributes to the strand of ex-post literature that tries to address the underlying endogeneity problem by employing an innovative method of comparative case studies, known as the synthetic control method. One of the key advantages of this approach is that SCM allows the unobserved confounders to vary with time, as opposed to other econometric methods that only allow for time-invariant unobservables. This is the first paper that employs SCM across a large number of trade agreements to answer a very relevant policy question in the current juncture: do trade agreements matter?

III. EMPIRICAL STRATEGY

> "Given that many policy interventions and events of interest in social science take place at an aggregate level (countries, regions, cities, etc.) and affect a small number of aggregate units, the potential applicability of synthetic control methods to comparative case studies is very large, especially in situation where traditional regression methods are not appropriate."
>
> Abadie, Diamond and Hainmueller (2010)

Synthetic control method (SCM) is a recent econometric tool for comparative studies developed by Abadie and Gardeazabal (2003) and later extended by Abadie, Diamond and Hainmueller (2010). SCM provides a systematic way of choosing comparison units in comparative studies as the choice of comparison unit is a data driven procedure. As highlighted in Hosny (2012), SCM is close in spirit to the matching estimators and difference-in-differences models, but rest on weaker identification assumptions that allows the effects of unobserved confounders to vary with time. SCMs create synthetic (artificial) control or comparison units based on their similarity to the treated country before the treatment actually takes place. This is done based on a weighted average of past observable covariates and past realizations of the outcome variables. The evolution of the actual outcome of the treated unit post-treatment is then compared against the outcome of the synthetic unit, and the difference is interpreted as the treatment effect.

Since the publication of Abadie, Diamond and Hainmuller (2010)'s seminal study of the effect of tobacco control programmes on tobacco consumption in California, the SCM approach has been very popular in micro and macro studies. The range of applications is very wide, from studies of natural disasters and political conflict through to social and economic policy interventions, as is the range of spatial scales, from whole countries to school districts. Abadie, Diamond and Hainmueller (2014) uses SCM to look at the impact of German reunification on economic growth in West Germany, while Lee (2011) uses the approach to evaluate the effect of inflation-targeting policies on inflation rates in emerging markets. A comprehensive list of studies using SCM during the period 2003-2015 can be found in Craig (2015).

While SCM has been a very popular approach in many micro and macro studies, there are relatively very few studies in international trade using this approach. Billmeier and Nannicini (2007) and Nannicini and Billmeier (2011) use SCM to look at trade openness and growth relationship. Hosny (2012) look at Algeria's trade with GAFTA countries.

A. SCM Methodology

This section summarizes the technical details underpinning the SCM methodology, following the exposition of Abadie, Diamond and Hainmueller (2010). The authors propose a simple model to provide the rationale for the use of SCM in comparative case study research. The model assumes that there are $J+1$ regions and that the first region is exposed to the intervention of interest, so that the remaining J regions are potential controls, or, as is known in statistical matching literature, "donor pool".

The authors make the following assumptions while building the model:

- Y_{it}^N is the outcome that would be observed for region i at time t in the absence of intervention, for units i=1,….,$J+1$, and time periods t=1,….,T.
- T_0 is the number of pre-intervention periods, with $1 \leq T_0 \leq T$.
- Y_{it}^1 is the outcome that would be observed for unit i at time t if unit i is exposed to the intervention in periods T_0+1 to T.
- The intervention has no effect on the outcome before the implementation period. Hence for $t \in [1, \dots, T_0]$ and all $i \in [1, \dots., N]$, $Y_{it}^1 = Y_{it}^N$.
- The observed outcome for unit i at time t is then $Y_{it} = Y_{it}^N + \alpha_{it} D_{it}$, where $\alpha_{it} = Y_{it}^1 - Y_{it}^N$ is the effect of the intervention for unit i at time t, and D_{it} is an indicator that takes value one if unit i is exposed to the intervention at time t, and value zero otherwise. Since only the first region is exposed to the intervention and only after period T_0:

$$D_{it} = \begin{cases} 1, & i = 1 \; and \; t > T_0 \\ 0, & otherwise \end{cases} \quad (1)$$

In this set-up, the aim is to estimate $(\alpha_{1T0+1}, + \cdots + \alpha_{1T})$.

$$\alpha_{1t} = Y_{1t}^1 - Y_{1t}^N = Y_{1t} - Y_{1t}^N \qquad for\ t > T_0 \quad (2)$$

Y_{1t} is observed. Hence, to estimate α_{1t}, the only variable that needs to be estimated is Y_{1t}^N. The idea of synthetic control is introduced here by Abadie, Diamond and Hainmueller (2010), in the process of estimating Y_{1t}^N. The authors assume that Y_{it}^N is given by a factor model:

$$Y_{it}^N = \delta_t + \theta_t Z_i + \lambda_t \mu_i + \varepsilon_{it}, \quad (3)$$

where δ_t is an unknown common factor with constant factor loadings across units, Z_i is a vector of observed covariates (not affected by the intervention), λ_t is a vector of unobserved common factors, μ_i is a vector of unknown factor loadings, and the error terms ε_{it} are unobserved transitory shocks at the region level with zero mean.

Abadie, Diamond and Hainmueller (2010) show that the equation (3) is a generalization of the difference-in-differences (fixed-effects) approach. However, a key difference is that the impact of unobservable common factor representing, in this example, region heterogeneity, λ_t, is allowed to vary with time. In difference-in-difference models, the unobserved heterogeneity is assumed to be constant across time, and can therefore be eliminated by taking differences.

As explained in Hosny (2012), Abadie, Diamond and Hainmueller (2010) further define the following $(J \times 1)$ vector of weights:

$$W = (w_2, \ldots\ldots, w_{J+1})\ such\ that\ w_j \geq 0\ and\ \sum w_j = 1\ for\ j = 2, \ldots.., J+1 \qquad (4)$$

Each particular value of W represents a potential synthetic control for the treated unit, that is, a particular weighted average of control units. The value of the outcome variable for each synthetic control indexed by W is:

$$\sum_{j=2}^{J+1} w_j Y_{jt} = \delta_t + \theta_t \sum_{j=2}^{J+1} w_j Z_j + \lambda_t \sum_{j=2}^{J+1} w_j \mu_j + \sum_{j=2}^{J+1} w_j \varepsilon_{jt} \qquad (5)$$

Suppose that there is an optimal vector $(w_2^*, \ldots.., w_{J+1}^*)$ such that

$$\sum_{j=2}^{J+1} w_j^*\, Y_{j1} = Y_{11}, \qquad \sum_{j=2}^{J+1} w_j^*\, Y_{j2} = Y_{12}, \quad \ldots, \quad \sum_{j=2}^{J+1} w_j^*\, Y_{jT0} = Y_{1T0} \quad and \quad \sum_{j=2}^{J+1} w_j^*\, Z_j = Z_1, \qquad (6)$$

then the following could be used as an estimator of the treatment effect,

$$\widehat{\alpha_{1t}} = Y_{1t} - \sum_{j=2}^{J+1} w_j^*\, Y_{jt} \qquad for\ \ t \in (T_0 + 1, \ldots, T) \qquad (7)$$

In practice, the synthetic control $(w_2^*, \ldots.., w_{J+1}^*)$ is selected so that equation (6) holds approximately. The vector $(w_2^*, \ldots.., w_{J+1}^*)$ is optimally chosen to minimize the following pseudo-distance:

$$\|X_1 - X_0 W\|\ such\ that\ w_j \geq 0\ and\ \sum w_j = 1\ for\ j = 2, \ldots, J+1, \qquad (8)$$

where X_1 represents a vector of pre-intervention characteristics of the treated region, while X_0 is a matrix containing the same pre-intervention variables of the control regions.

In simple words, the SCM essentially uses a weighted average of the outcome of potential control units to estimate the counterfactual outcome of the treated unit. The weights are chosen in such a way that, on average, the estimated pre-intervention outcome and observed characteristics of the synthetic control are very close to those actually observed in the treated unit. One of the novelties of this approach is that, unlike difference-in-difference models, it allows the impact of unobservable confounding factors to be time-invariant.

While SCM approach has been extremely popular in recent studies to address endogeneity concerns, some of the disadvantages should be borne in mind, well summarized in Craig (2015). First, since the method relies on comparisons between a real region and a synthetic control unit, standard methods of statistical inference are inappropriate. Instead, Abadie, Diamond and Hainmueller (2010) propose to use placebo or 'falsification' tests described in Section IV. Hence, inference is less formal compared to other econometric methods and is based on weaker assumptions. Second, the effect of the intervention can only be estimated accurately if there were no other events that affect only the intervention area, such as additional policy changes. Third, SCM assumes that the intervention should affect outcomes only in the intervention area. If the other areas comprising the synthetic unit are also affected, the impact could be potentially underestimated. Fourth, a good fit is needed between the pre-intervention trends in the intervention area and the synthetic control. Finally, for the method to work well, some weighted combinations of regions in the donor or control pool must be similar to the intervention area. If the treated region is at the extreme end of the range of observed characteristics, it might be difficult to find an appropriate synthetic.

B. Application of SCM to Bilateral Trade Pairs

This section discusses how SCM can be used to determine the impact of trade agreements. First, the trade agreements – the treatment in SCM approach – are discussed. Second, gravity model is used to determine the covariates that are used to form the synthetic pair. Finally, an illustrative example is given to clarify the concept.

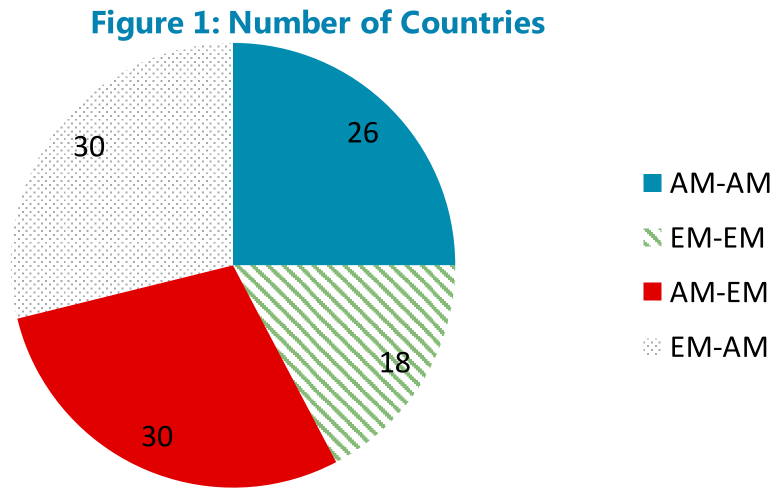

The Trade Agreements

This paper looks at the impact of regional trade agreements in the period 1983-1995[2]. The regional trade agreements are taken from Head, Mayer and Ries (2010). Each agreement between two countries, for example Country A and Country B, would translate into two pairs, one representing the exports from Country A to Country B, the other representing the exports from Country B to Country A. Using this logic, there are 104 bilateral country pairs in the period considered. Figure 1 shows the income dynamics of the pairs studied. The notation EM-AM corresponds to an emerging country exporting to an advanced economy. Of the 104 pairs studied, there are 30 EM-AM, 30 AM-EM, 26 AM-AM and 18 EM-EM pairs.

From Gravity Model to SCM

This paper employs SCM approach to understand bilateral exports of countries for the ten years following a trade agreement. The trade agreement is the 'treatment' undergone while the country pair with the trade agreement can be regarded as the 'treated' unit. To construct the synthetic that would represent the counterfactual, the donor pool (or the control group) should naturally not include the same treatment. Consequently, the donor pool for each treated unit excludes all the country pairs in the sample that had a trade agreement in the same year. In addition, the donor pool excludes all other trade agreements that the exporting

[2] For robustness check, the paper also looks at a broader sample comprising 216 country pairs in the period 1973-2001.

country has with any other country in the entire sample. For few cases where the program is not successful in determining synthetic using the aforementioned donor pool, an alternative donor pool is used that excludes only trade agreements with the country that undergoes the trade agreement.

As discussed in the previous section, the synthetic variable is constructed based on a weighted average of past observable covariates and past realizations of the outcome variable. For our purpose, in order to obtain the covariates, we need a model that can be useful in explaining bilateral trade between two country pairs. We can then use the determinants of trade (or regressors/covariates) from that model and the exports between two country pairs (or outcome variable) to create the synthetic counterfactual outcome. For the regressors or covariates of the SCM approach, we resort to the gravity model that has been widely used in the last four decades in empirical analysis to understand patterns and potential of trade. Feenstra, Markusen and Rose (2001) argue that simple gravity equation explains a great deal about the data on bilateral trade flows and is consistent with several theoretical models of trade.

The gravity model is essentially a simple equation to predict the trade on a bilateral basis between any two countries. The name is derived from the similarity to the law of gravity in physics, with the attraction (trade) depending on mass (economic size) and distance. Recent studies have augmented the simple model to include dummy variables incorporating some characteristics common between two countries that may be useful determinants of bilateral flows. As discussed in Head, Mayer and Ries (2010), all the well-known empirical and theoretical formulations of the gravity equation can be represented in the following equation for the value of x_{ijt}, the exports from exporting country *i* to importing country *j* in year *t*:

$$x_{ijt} = G_t M_{it}^{ex} M_{jt}^{im} \emptyset_{ijt} \qquad (9)$$

where M_{it}^{ex} and M_{jt}^{im} are indexes of the attributes of exporter *i* and importer *j* in a specific year. As highlighted in UNCTAD/WTO (2012), M_{it}^{ex} comprises export-specific factors (such as the exporter's GDP) that represent the total amount exporters are willing to supply and M_{jt}^{im} denotes all importer-specific factors that make up the total importer's demand (such as the importing country's GDP). G_t is a common year-specific factor determining trade that does not depend on i or j such as the level of world liberalization. Finally, variation in bilateral trade intensity enters through $\emptyset_{ijt}$, a variable that represents the ease of exporter i to access of market j (that is, the inverse of bilateral trade costs). In literature, M_{it}^{ex} and M_{jt}^{im} are referred to as monadic effects while $\emptyset_{ijt}$ as the dyadic effect. In practice, the gravity equation relates the monetary value of trade between two countries to their respective GDPs, composite terms measuring barriers and incentives to trade between them, and terms measuring barriers to trade between each of them and the rest of the world. UNCTAD/WTO (2012) has a detailed exposition on this.

In line with literature on the determinants of trade, the paper uses the following variables for covariates in the SCM approach:

- Distance between the bilateral pairs
- GDP of each country in the bilateral pair
- GDP per capita of each country in the bilateral pair
- Population of each country in the bilateral pair
- Bilateral Real Exchange Rate
- Remoteness of each country in the bilateral pair, proxy for multilateral trade resistance (MTR) term (remoteness due to physical distance and/or policy).
- Dummy variables:
  - Colonial history = 1 if pair ever in colonial relationship
  - Col to = 1 if export from hegemon to colony
  - Col from = 1 if export from colony to hegemon
  - Contig = 1 for contiguity
  - Comleg = 1 for common legal origins
  - Comcur = 1 for common currency
  - Common language = 1 for common official language
- Exports, lagged by three years

The data comes mostly from the underlying gravity dataset used in Head, Mayer and Ries (2010), available at http://econ.sciences-po.fr/thierry-mayer/data. The authors note that the trade data in the dataset comes from IMF DOTS database. The trade agreements are taken from three sources: Baier and Bergstrand (2007), supplemented with information from WTO and qualitative information contained in Frenkel (1997). The gravity controls are taken mainly from World Bank's Development Indicators, national sources, and Angus Maddison dataset. The authors also use data from Andrei Shleifer and Glick at http://post.economics.harvard.edu/faculty/shleifer/Data/qgov_web.xls and Glick and Rose

(2002) for common legal origins and common currency, respectively. Bilateral distances and common (official) language come from the CEPII distance database (http://www.cepii.fr/anglaisgraph/bdd/distances.htm). Outside the dataset by Head, Mayer and Ries (2010), bilateral real exchange rate is calculated using data from World Development Indicators. The entire sample contains 38124 observations.

All the variables are self-explanatory. However, multilateral trade resistance (MTR) warrants some discussion. This term arises from the seminal work of Anderson and van Wincoop (2003), who show that bilateral trade is determined by related trade costs. The propensity of country *j* to import from country *i* is determined by country's trade cost toward *i* relative to its overall resistance to imports (weighted average trade costs) and to the average "resistance" facing exporters in country *i*; not simply by the absolute trade costs between countries *i* and *j*. Following the practice of a lot of studies in literature (UNCTAD/WTO, 2012), this paper incorporates MTR terms for exporting and importing countries by including a proxy for these indexes called "remoteness":

$$Rem_i = \sum_j \frac{dist_{ij}}{GDP_j / GDP_w} \qquad (10)$$

This formula measures a country's average weighted distance from its trading partners (Head, 2003), where weights are the partner countries' shares of world GDP (denoted by GDP_w).

An Illustrative Example

This section gives an illustrative example to understand how SCM works, particularly in the context of bilateral trade agreements. The treated unit is Belgium exports to Spain (B-S) and the treatment is the EC (12) enlargement in 1986. In other words, one is interested in understanding the impact of the regional trade agreement on Belgium exports to Spain. The challenge is to produce a 'synthetic' B-S pair using a control group of bilateral pairs that have not been exposed to the same trade agreement. The synthetic pair is the weighted average of potential 'control' bilateral trade pairs, with weights chosen so that the resulting synthetic pair best reproduces the values of bilateral exports from Belgium to Spain before

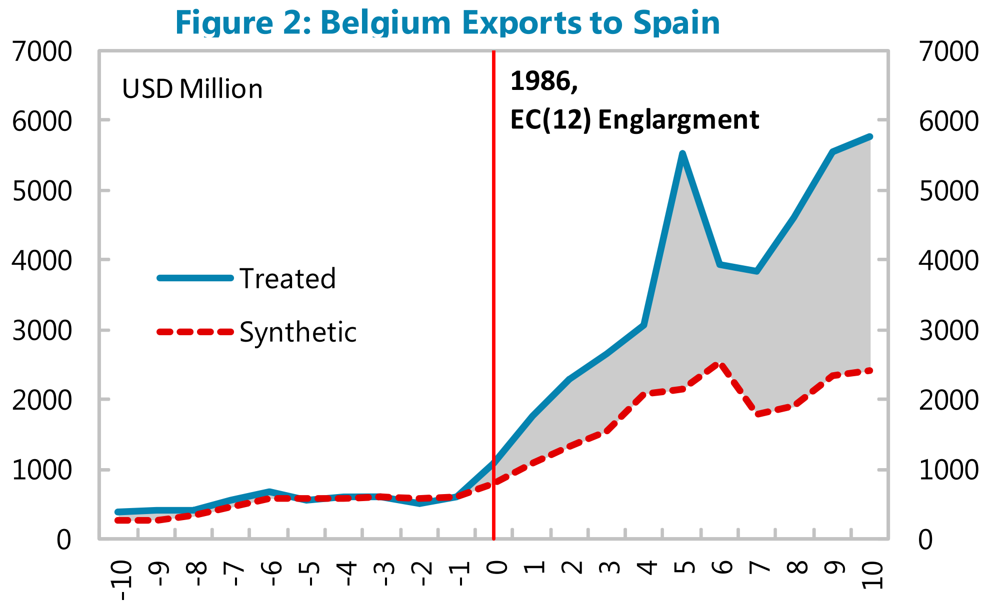

the trade agreement. The synthetic B-S pair is thus meant to reproduce the Belgium exports to Spain that would have been observed in the absence of the trade agreement.

Figure 2 shows the results, with both treated unit – the actual B-S pair – and it's synthetic. The close match between the treated and the synthetic unit before the treatment shows that, in this case, the program was successful in finding a good synthetic pair for B-S. Consequently, the shaded area between the treated and the synthetic unit represents the Belgium's export gains vis-à-vis Spain due to the EC enlargement. Note that, as highlighted by Abadie, Diamond and Hainmueller (2010), in practice, interventions may have an impact prior to their implementation (e.g., via anticipation effects). This explains the divergence between the treated and the synthetic unit one year prior intervention.

IV. Results

This section discusses the results in two groups. The first part discusses the results on export gains due to trade agreement with the counterpart with which the country had a trade agreement. This part also includes results on placebo tests, usually done for this sort of comparison case analysis. The second part looks at the issue of trade diversion, both import and export diversion.

Figure 3: Exports, Average of 104 Country Pairs

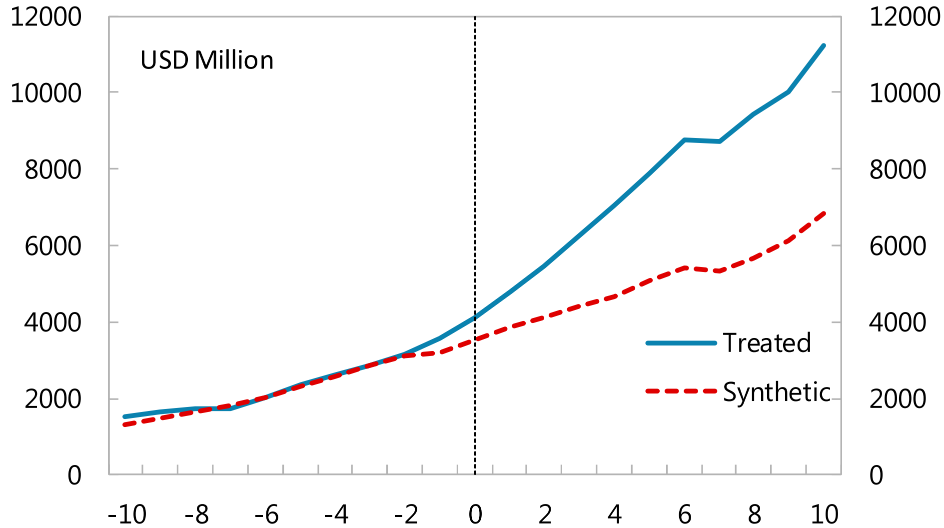

A. Trade Creation

As discussed in the previous section, the synthetic control method is employed to 104 bilateral country pairs that engaged in a trade agreement in the period 1983-1995. Figure 3 shows the results for the average of the 104 country pairs, while figures 4 and 5 look specifically at NAFTA trade agreement and 1986 EC(12) Enlargement[3]. The outcome variable, gross exports, is plotted against the time to trade agreements, with t=0 referring to the period when the agreement was enacted, t=-1,-2,…,-10 showing the outcome variable until ten years before the trade agreement, and t=+1, +2, …, +10 showing the outcome variable until ten years after the trade agreement. The blue line shows the average gross exports of countries undergoing the trade agreement, while the red dotted line shows their synthetic counterparts. In order to get a sense of the fit between the treated and synthetic before intervention, the normalized root-mean-square deviation (NRMSD)[4] or error is

[3] For NAFTA, bilateral country pairs corresponding to Canada and USA trade has t=0 in 1989, reflecting the Canada-United States Free Trade Agreement that was superseded by NAFTA. Taking t=0 in 1994 for these pairs yields similar overall results: the average export gains for NAFTA is 72 percentage points in ten years, slightly lower than the baseline result of 79 percentage points reported in the text.

[4] The root-mean-square deviation (RMSD) or root-mean-square-error (RMSE) is a frequently used measure of the differences between values predicted by a model and the values actually observed. RMSD is the square root of the mean of the squares of deviations. Normalizing the RMSD facilitates the comparison between different

Figure 4: Exports, Average of Pairs in NAFTA

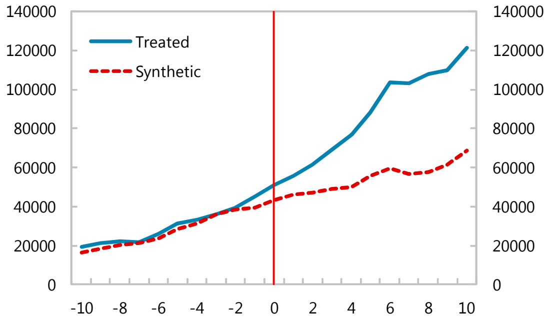

Figure 5: Exports, Average of Pairs in EC(12)

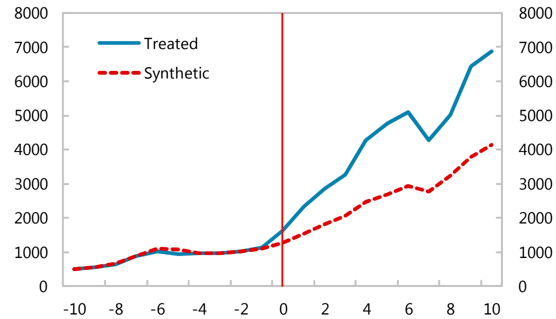

computed for the ten years prior to intervention. The average NRMSD across the 104 pairs is 0.56, while the median of the same measure is 0.31. In other words, both figure 3 and the low NRMSD suggests good fit between treated and synthetic before intervention.

The divergence between the treated and synthetic lines in all the figures post trade agreement indicates substantial gains from trade[5]. The average gross exports of countries with trade agreement was 80 percentage points higher over the next ten years on a cumulative basis, compared to the average synthetic counterparts (figure 6). This translates into annual average export growth of 3.8 percentage points higher than the average export

Figure 6: Export Growth of Treated Over Ten Years, Relative to Synthetic (Cumulative, Percentage Points)

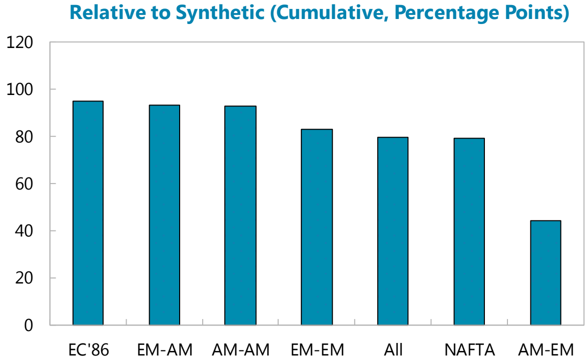

Figure 7: Annual Export Growth of Treated, Relative to Synthetic (Percentage Points)

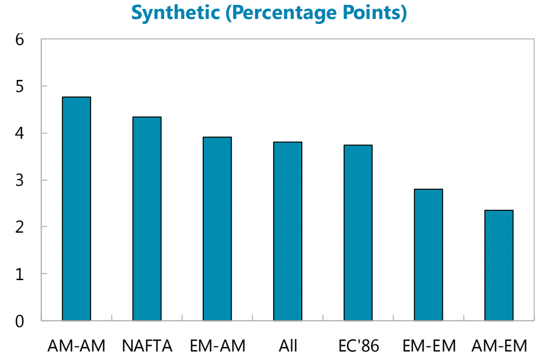

units. The normalized root-mean-square deviation (NRMSD) is computed by dividing the RMSD of each country pair by the average exports over ten years prior to intervention.

[5] However, it must be noted that not all the countries gained from trade. Of the 104 country pairs considered, 80 countries had treated higher than synthetic after ten years following the trade agreement, while 24 countries had treated equal or less than synthetic.

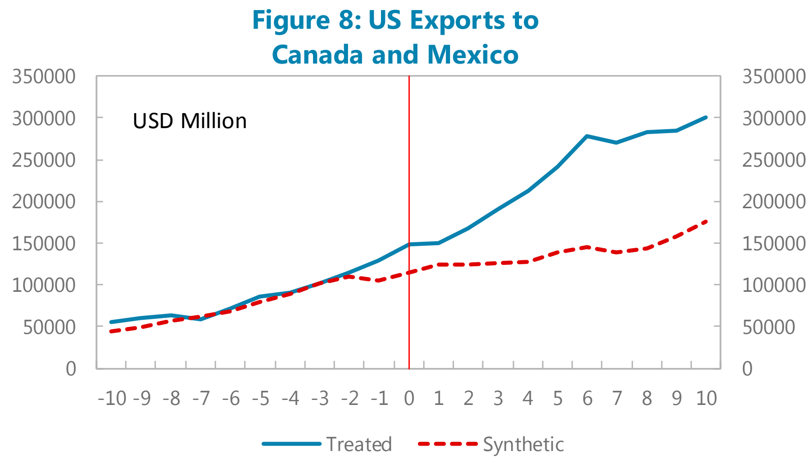

growth in the synthetic (figure 7)[6]. Looking across income groups, the gains from trade is the highest for EM-AM pairs (an emerging market exporting to advanced economy) with 93 percentage points increase over ten years on a cumulative basis relative to the synthetic, translating into around 3.9 percentage points higher annual exports growth. The gains from AM-AM pairs (an advanced market exporting to advanced economy) are also high. The gains are relatively less, compared to the average, for advanced economies exporting to emerging markets.

Looking at some specific trade agreements, the 1986 EC Enlargement generated 95 percentage points higher export growth in 10 years, while the average annual export growth was 3.7 percentage points higher. The exports due to NAFTA increased by 79 percentage points higher than synthetic over ten years, and the average annual export growth was 4.3 percentage points higher. The export gains of the US due to NAFTA was also substantial, evident from a higher trajectory of treated compared to the synthetic in figure 8, showing the US total exports to Canada and Mexico. For the US, this translates to export growth of around 50 percentage points higher over ten years, with annual average export growth 3.0 percentage points higher than the synthetic counterpart.

[6] Figures 6 and 7 are based on the assumption that the gains of trade start at t=0, the year when the trade agreement is enacted. However, in reality, some of the effects of trade agreements could occur before the promulgation due to the anticipation effect. If the trade gains are calculated from the period when the treated line diverges from the synthetic counterpart, (instead of the time when the agreement was enacted) until the ten years following the trade agreements, the average export growth for the 104 country pairs is 132 percentage points higher than the average synthetic. For NAFTA, the export growth is 128 percentage points higher, while for the 1986 EC Enlargement, it is 234 percentage points higher.

For robustness check, the paper also looks at a broader sample comprising 216 country pairs in the period 1973-2001. This was not used as the baseline as there were data gaps. With that caveat in mind, the results using the longer sample shows that the average export gains due to trade agreements is 90 percentage points, with an annual export growth of 3.8 percentage points. The results are thus stronger than using the shorter sample.

These results add to the existing literature in few dimensions. This study falls under a small group of literature that shows that, once endogeneity issues are addressed, trade gains can be substantial. As noted in Baier and Bergstrand (2007), one of the puzzles in trade literature is that one would expect from the sheer number of trade agreements that they matter. However, academic literature has shown little impact of trade agreements in boosting trade. Compared to Baier and Bergstrand (2007, 2009) – one of the few studies showing large trade gains – who find that trade gains are 100 percentage points higher over ten years, the results are lower when export gains are calculated using t=0. However, the results are higher when the anticipation effect is accounted for. For NAFTA, one key finding is that the treated unit is larger than the synthetic unit for all the possible bilateral pairs involving the US, Canada and Mexico, indicating that the exports have increased for all the participating countries in NAFTA. Literature on NAFTA usually does not find trade gains across all the bilateral country pairs. Gould (1998) show that only the US exports to Mexico is statistically significant when looking at the impact of NAFTA. In addition, the results for NAFTA are, broadly speaking, higher than the estimates suggested by key ex-post literature (Krueger, 1999; Wall, 2003). Krueger (1999) studies aggregate and micro data on trade between the US, Canada and Mexico to conclude that the impact of NAFTA over its first three years was around 3 percent increase in trade and the results were not statistically significant.

Another dimension where this paper shed light is the differentiation across income groups, which literature usually does not address. These results show that the gains of trade could depend on the income groups of the countries involved. This could be due to couple of factors. First, the analysis here as well as literature implicitly assumes that all trade agreements are same. However, the nature of trade agreements could differ. For example, an advanced economy’s trade agreement with another advanced economy could be deeper involving aspects like regulatory cooperation, competition, etc. Second, for emerging markets, trade agreements with another advanced economy could potentially expose them to large markets for their products, resulting in more trade growth.

Placebo (or Falsification) Tests

This section assesses whether the effect estimated by the synthetic control method for a country pair affected by the trade agreement is large relative to the effect estimated for a country pair chosen at random. For SCM studies, Abadie, Diamond and Hainmueller (2010) propose exact inferential techniques, akin to permutation tests, to perform inference. The idea is to apply the synthetic control method to the controls in the sample and determine if the effect estimated by the synthetic control for the country pair with trade agreement (that is, affected by intervention) is large relative to the effect estimated for a country pair chosen at random that did not go through the trade agreement.

The process of performing placebo tests is as follows: ten treated units are randomly chosen from the sample. Let A be the exporting country of a treated unit chosen randomly. The idea is then to randomly select five country pairs showing the exports of A to a country that does not have a trade agreement with A. These five country pairs are known as placebos. This process of choosing ten random treated units and five control units per treated unit should thus generate fifty control units (placebos). SCM is then applied to these control units. The evolution of treated unit relative to synthetic unit is then compared to the evolution of the placebo units relative to their synthetics.

Figure 9 shows the results of this exercise, where the bolded black line represents treated relative to synthetic and the other lines in each chart represent the evolution of the placebo relative to synthetic. Of the ten treated units randomly chosen, there are eight cases where treated is greater than its synthetic pair. There are two cases, the last two charts in light blue, where treated is less than synthetic. Of the eight cases where the treated is higher than synthetic, the placebo tests show that there are seven cases were the treated relative to synthetic is higher than placebo relative to synthetic. The placebo test gives one failure, the chart in grey. There is only one case where the treated is higher than synthetic and treated relative to synthetic is lower than placebo relative to synthetic. This indicates that the placebo test has been successful and the inferential technique, as implied by placebo tests using the approach suggested by Abadie, Diamond and Hainmueller (2010), is strong.

Figure 9: Placebo Tests: Treated (or Placebo) Relative to Synthetic

Black line is treated, others placebo (USD Million)

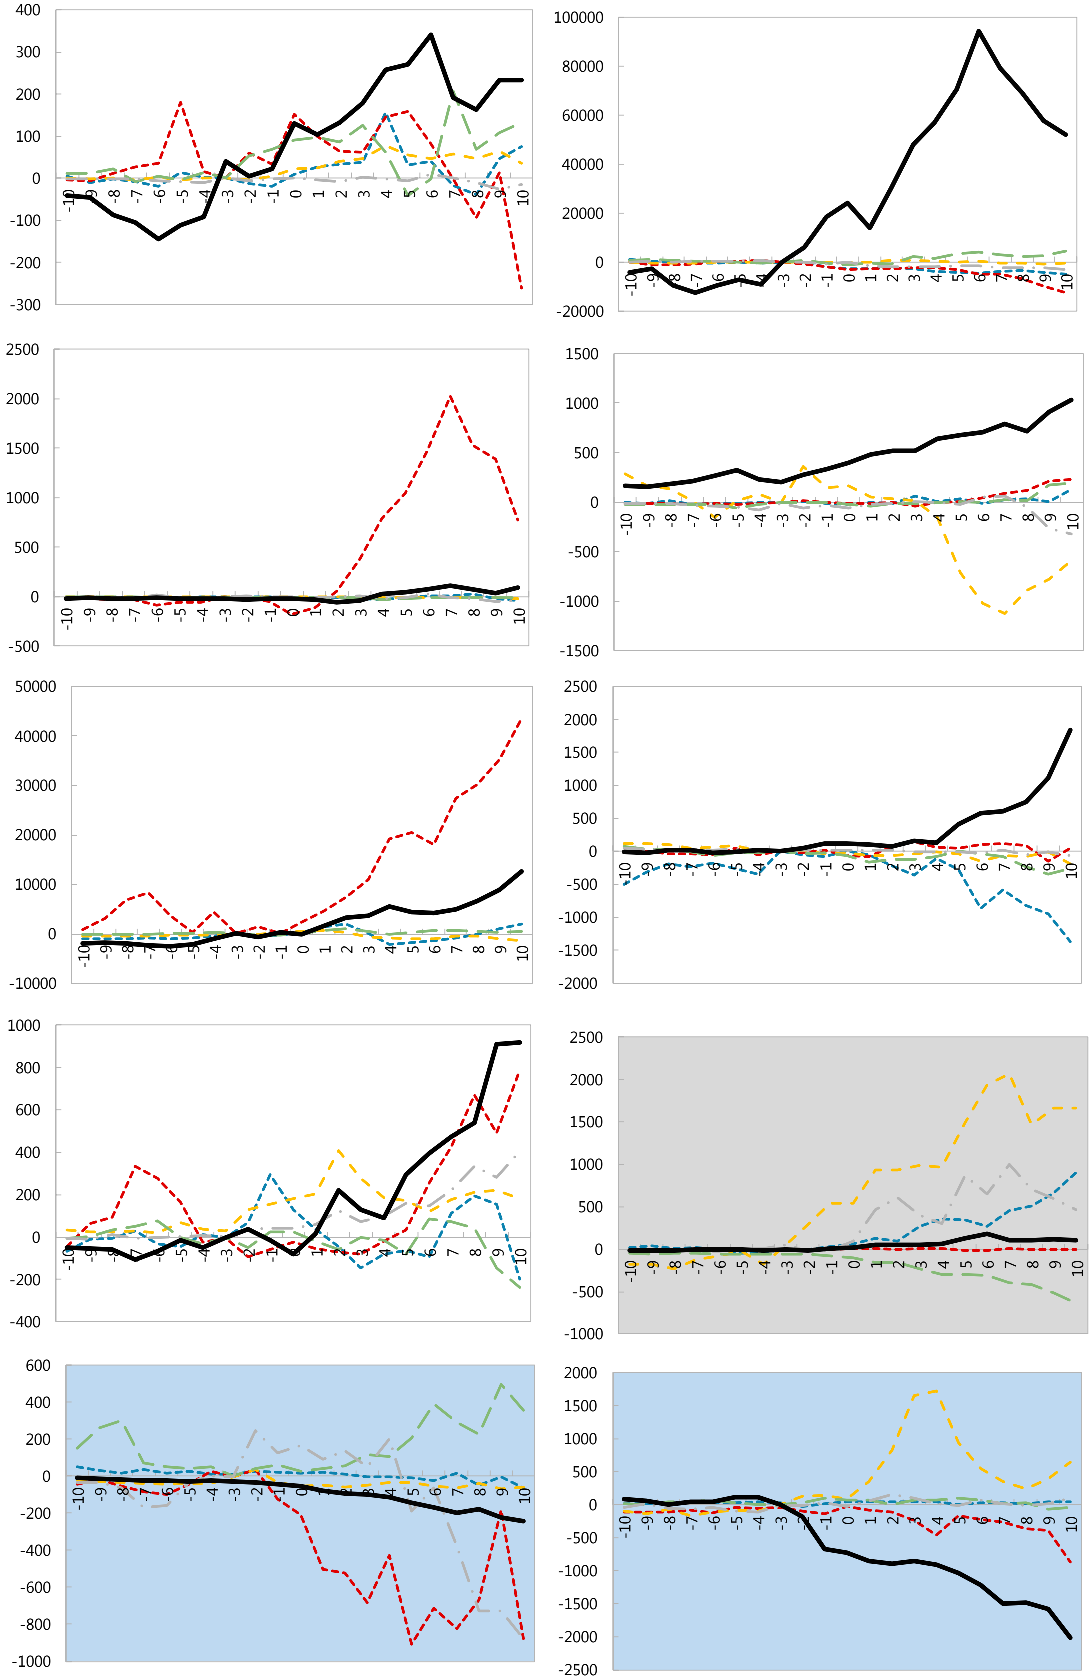

B. Trade Diversion

This section determines if the trade creation, suggested by the positive export gains in the previous section, occurs at the expense of other countries that are not part of trade agreements. Is there evidence of trade diversion? This section exploits the underlying logic of SCM to look at trade diversion in a novel manner. For each country that has engaged in a trade agreement, the idea is to apply SCM to the pair representing the country's trade with its top trading partners that are outside the trade agreement. The synthetic unit in this scenario would represent the counterfactual of how trade would have evolved with the third country under the absence of trade agreement. While constructing the synthetic, the control group excludes the same country pairs as in the analysis for trade creation. Like the previous section, there were few cases where the control group was extended to include the bilateral pairs with trade agreements as well, mainly due to similar reasons.

Import Diversion

Let countries A and B be engaged in a trade agreement. The idea here is to look at country A's top importer outside the trade agreement. Let C be that country. Here, applying SCM would answer the question: what would have happened to country A's imports from country C if country A did not have a trade agreement with country B? Using this idea for the

Figure 10: Imports from Countries Not in Trade Agreement

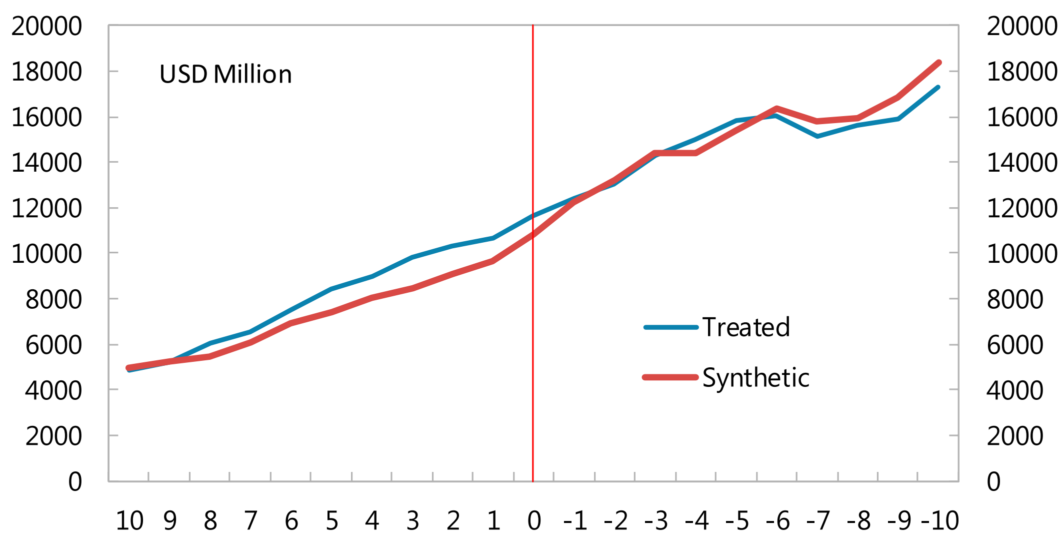

Figure 11: Import Growth of Treated Over Ten Years, Relative to Synthetic (Cumulative, Percentage Points)

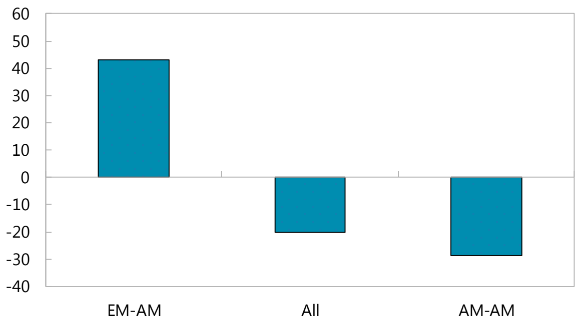

Figure 12: Average Import Growth of Treated, Relative to Synthetic (percentage points)

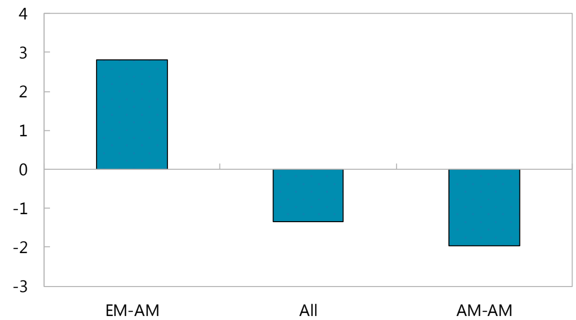

104 country pairs results in 43 country pairs: 31 AM-AM (advanced economy importing from another advanced economy) and 12 EM-AM (emerging economy importing from an advanced economy). Figure 10 shows the evolution of the average treated[7] (or in this context, the pair that did not go through treatment) and synthetic units, and figures 11 and 12 summarize the results. The fact that the synthetic unit is slightly higher than the treated unit suggests that the imports from the countries would have been slightly higher if there were no trade agreement. In other words, there is some suggestive evidence for import diversion. The synthetic unit is 20 percentage points higher than the treated unit after ten years (on a cumulative basis). Looking at income groups, there seems to be import diversion from AM-AM pairs but EM-AM pairs show positive results. However, the results should be interpreted with a pinch of salt since the fit – the match between treated and synthetic – before intervention is not the best, particularly compared to the fit witnessed when analyzing trade creation and export diversion (next section).

Export Diversion

The same idea is used look at export diversion. Figure 13 shows the evolution of the average outcome and the average synthetic, while figures 14 and 15 summarize the results. The treated unit is higher than its synthetic counterpart, indicating that there is no evidence of export diversion. Quite the contrary, the results indicate that exports to third parties not part

[7] Note that, in figures and text, 'treated' unit is used to match the terminology of SCM. However, in the section on trade diversion, the treated unit has not undergone any treatment. In other words, the treated unit does not have a trade agreement.

of trade agreement could get a small boost due to the trade agreement. The export growth to third party is 8 percentage points higher over ten years due to the trade agreement. For emerging markets exporting to advanced economies, the export gains are around 19 percentage points over ten years while the average annual export growth is 0.9 percentage points higher.

The results of import and export diversion should be interpreted with the caveat that this is based on a narrow analysis looking at only the country's top importer/exporter outside the trade agreement. While this gives some important clues regarding trade diversion, future research could expand this to include SCM for more country pairs that are not part of the trade agreement.

Figure 13: Exports to Countries Not in Trade Agreement

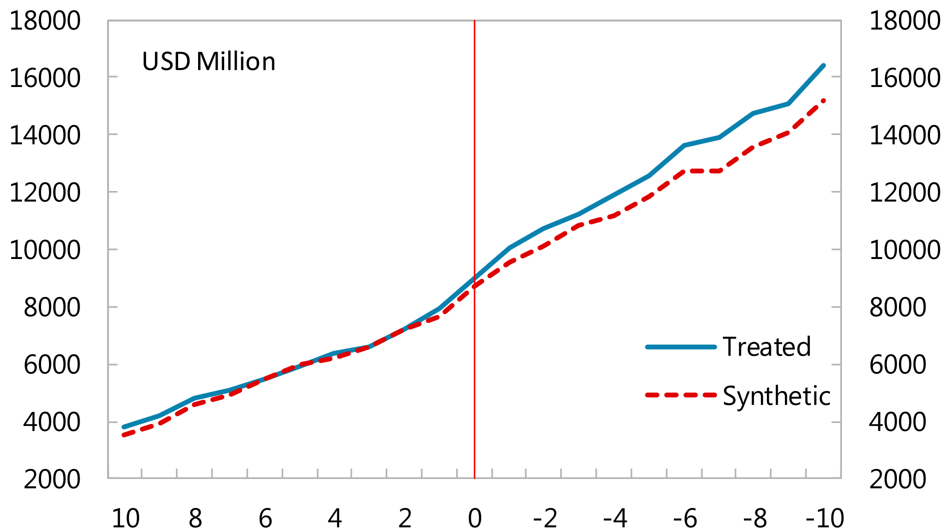

Figure 14: Export Growth of Treated Over Ten Years, Relative to Synthetic (Cumulative, Percentage Points)

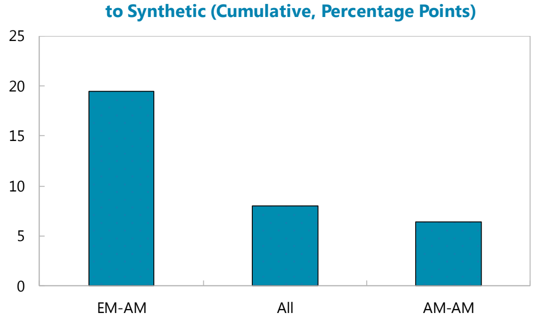

Figure 15: Annual Export Growth of Treated, Relative to Synthetic (Percentage Points)

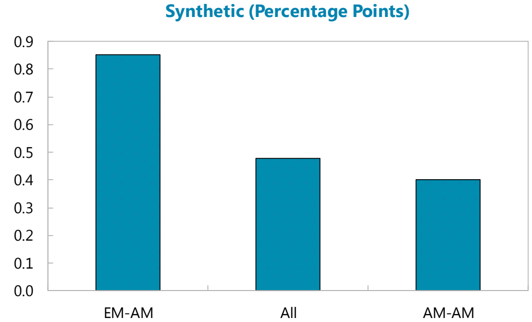

V. Concluding Thoughts

Against the backdrop of renewed interest on the impact of trade agreements due to TPP, this paper employs a novel approach – synthetic control methods – to determine how trade agreements have influenced trade. The results show that trade agreements can substantially boost exports, by an average of 80 percentage points over ten years on a cumulative basis with an annual average export growth 3.8 percentage points higher. The export gains are higher for emerging markets in trade agreements with advanced markets. Looking at specific trade agreements, the results show that all the bilateral country pairs in NAFTA have substantially increased exports due to NAFTA, resulting in 79 percentage points higher exports over ten years, with an average export growth of 4.3 percentage points. The paper also looks at the issue of trade diversion by determining what happens to a country's top exporter and top importer that are outside the trade agreement. With the caveat that the fit is not the best while looking at top importers, the results give suggestive evidence of import diversion, but not export diversion. However, the magnitude of import diversion is small, around 20 percentage points over ten years, which would be more than offset by the trade creation. Hence, putting everything together, the results would indicate substantial trade creation.

The SCM is currently very popular in micro and macro studies since it can allow for unobserved confounders to vary with time, as opposed to traditional econometric methods that allows for only time-invariant unobservables. The simple and intuitive data-driven approach of determining the control group also attracts researchers to this approach. However, some of the limitations of the approach, discussed in detail in Section III, should be borne in mind. Since the method relies on comparisons between a real area and a synthetic control area, standard methods of statistical inference are inappropriate and one has to resort to placebo tests as an inferential technique. The method also depends on the assumption that no other intervention took place at the same time, and the results could thus be contaminated if there were other non-trade policy changes that could affect trade developments. Another assumption that may not always hold is that the intervention should affect outcomes only in the intervention area.

The scope of the analysis should also be taken into consideration when interpreting the results. The focus of interest or the outcome variable in this paper is exports. While thinking about the broader issue of trade gains for the economy, there are many other factors that need to be taken into consideration, including welfare gains, labor adjustments, impact on inequality, costs of trade liberalization to the society. Fajgelbaum and Khandelwal (2016) develop a methodology to measure the unequal gains from trade across consumers within countries and find that trade typically favors the poor, who concentrate spending in more traded sectors. On the other hand, Autor, Dorn, Hanson and Song (2014) analyze the effect of trade exposure on earnings and employment of U.S. workers from 1992 through 2007. The authors find that import shocks impose substantial labor adjustment costs that are highly unevenly distributed across workers according to their skill levels and conditions of employment in the pre-shock period. Future studies could use SCM approach to look at some of these outcome variables.

One might expect that the sheer number of trade agreements in the past decades would indicate that such agreements can boost trade. However, there is very little empirical support to this claim (Baier and Bergstrand, 2007). The paper belongs to the small group of literature that shows that, after accounting for endogeneity, trade agreements matter! This is particularly relevant for policy making in the current context of trade slowdown witnessed in data (Constantinescu, Mattoo and Ruta, 2015), and more work in employing SCM in trade relevant issues would be welcome. Future studies could expand the SCM analysis to understand how different sectors react to trade agreements. In addition, the analysis on trade diversion could be expanded to include more countries and sectors to understand which sectors suffer most if excluded from trade agreement. Finally, the impact of trade agreements in the formation of global value chains could also be studied exploiting this novel approach.

VI. References

Abadie, A., A. Diamond, and J. Hainmueller (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.

Abadie, A., A. Diamond, and J. Hainmueller (2014). “Comparative Politics and the Synthetic Control Method”, American Journal of Political Science, 59 (2).

Abadie, A. and J. Gardeazabal (2003). “The economic costs of conflict: A case study of the Basque Country”, American Economic Review, 93(10, 113–132.

Aichele, R., and G. Felbermayr (2016). "The Trans-Pacific Partnership Deal (TPP): What Are the Economic Consequences for In- and Outsiders?" Ifo Institute – Leibniz Institute for Economic Research at the University of Munich, 16(4), 53–64.

Anderson, J. E. and E. van Wincoop (2003), "Gravity with gravitas: a solution to the border puzzle", American Economic Review 93, 170–92.

Baier, S.L. and J.H. Bergstrand (2004). "Economic determinants of free trade agreements", Journal of International Economics, 64, 29–63.

Baier, S.L. and J.H. Bergstrand (2007). "Do free trade agreements actually increase members' international trade?", Journal of International Economics, 71, 72–95.

Baier, S.L. and J.H. Bergstrand (2009). "Estimating the effects of free trade agreements on international trade flows using matching econometrics", Journal of International Economics, 77, 63–76.

Bergstrand, J. H. (1985). "The gravity equation in international trade: some microeconomic foundations and empirical evidence", Review of Economics and Statistics 67 (3), 474–481.

Billmeier A., and T. Nannicini (2007). "Trade openness and growth: pursuing empirical glasnost", IMF Working Paper, WP/07/156, International Monetary Fund, Washington, DC.

Carrere, C. (2006). "Revisiting the effects of regional trade agreements on trade flows with proper specification of the gravity model", European Economic Review, 50, 223–247.

Constantinescu, C., A. Mattoo, and M. Ruta (2015). "The Global Trade Slowdown: Cyclical or Structural?", IMF Working Paper, WP/15/6, International Monetary Fund, Washington, DC.

Corcos, G., M. del Gatto, G. Mion, and G. Ottaviano (2012). Productivity and Firm Selection: Quantifying the New Gains from Trade, Economic Journal, 122, 754-798.

Craig, P. (2015). "Synthetic controls: a new approach to evaluating interventions", Working Paper, What Works Scotland.

Ciuriak, D., and J. Xiao (2014). "Impact of Taiwan's TPP Accession on the United States", available at SSRN: http://ssrn.com/abstract=2547170.

Fajgelbaum, P. D. and A.K. Khandelwal (2016). "Measuring the Unequal Gains from Trade", The Quarterly Journal of Economics (*forthcoming)*.

Feenstra R. C., J. R. Markusen, and A. K. Rose (2001). Using the gravity equation to differentiate among alternative theories of trade", Canadian Journal of Economics, 34(2), 430–447.

Frankel, J. (1997). "Regional Trading Blocs", Institute for International Economics, Washington, DC.

Ghosh, S. and S. Yamarik (2004). “Are regional trading arrangements trade creating? An application of extreme bounds analysis”, Journal of International Economics 63 (2), 369–395.

Glick, R., and A. Rose (2002). “Does a currency union affect trade? The time-series evidence”, European Economic Review, 46 (6), 1125–1151.

Frankel, J., E. Stein, and S. Wei (1995). “Trading blocs and the Americas: The natural, the unnatural, and the super-natural”, Journal of Development Economics, 47 (1), 61–95. Available at www.sciencedirect.com/science/article/pii/0304387895000054.

Head, K. (2003). “Gravity for beginners”, mimeo, University of British Columbia.

Head, K., T. Mayer, and J. Reis (2010), “The erosion of colonial trade linkages after independence”, Journal of International Economics, 81, 1–14.

Hosny, A. S. (2012). “Algeria’s Trade with GAFTA Countries: A Synthetic Control Approach”, Transition Studies Review, 19, 35–42.

Hufbauer, G. C. and J. Schott (2005). “NAFTA Revisited, Achievements and Challenges”, Institute for International Economics, Washington, DC.

Kehoe, T. J. (2003). “An Evaluation of the Performance of Applied General Equilibrium Models of the Impact of NAFTA”, Staff Report 320, Federal Reserve Bank of Minneapolis.

Lee, W. S. (2011), “Comparative Case Studies of the Effects of Inflation Targeting in Emerging Economies”, Oxford Economic Papers, 63 (2), 375–397.

Lipsey, R. G. (1960). “The Theory of Customs Unions: A General Survey”, Economic Journal, 70 (289), 496–513.

Magee, C. (2003). “Endogenous preferential trade agreements: an empirical analysis. Contributions to Economic Analysis and Policy”, vol. 2(1), Berkeley Electronic Press.

Mansfield, E.D. and E. Reinhardt (2003). “Multilateral Determinants of Regionalism: The Effects of GATT/WTO on the Formation of Preferential Trade Agreements”, International Organization, 57 (4), 829–862.

Mansfield, E.D. and E. Reinhardt (2008). “International Institutions and the Volatility of International Trade”, International Organization, 62 (4), 621–652.

Petri, P., and M. Plummer (2016). “The Economic Effects of the Trans-Pacific Partnership: New Estimates”, Peterson Institute for International Economics, Working Paper No. 16–2.

Nannicini T., and A. Billmeier (2011). “Economies in transition: How important is trade openness for growth?” Oxford Bulletin of Economics and Statistics, 73(3), 287–314.

Tinbergen, J. (1962). “Shaping the World Economy”, The Twentieth Century Fund, New York.

UNCTAD Virtual Institute and World Trade Organization (2012). “A Practical Guide to Trade Policy Analysis”, Chapter 3, http://vi.unctad.org/tpa/index.html.

Viner, J. (1950). "The Customs Union Issue", New York: Carnegie Endowment for International Peace.