The Macroeconomic Effects of the Buy American Act

Miguel Acosta
University of Wisconsin–Madison

Lydia Cox*
University of Wisconsin–Madison and NBER

July 6, 2025

Abstract

We ask whether Buy American restrictions on government procurement lead to stronger macroeconomic effects of government spending. We propose a methodology based on a comprehensive understanding of Buy American regulations to identify government spending that is more and less constrained by domestic content restrictions. We validate our methodology in a novel firm-level dataset, and then apply it to federal spending at the state level. We show that government dollars that are not constrained by domestic content restrictions have a higher cross-sectional fiscal multiplier than constrained contract dollars. We rationalize our findings in the context of a New Keynesian model of U.S. regions. Though Buy-American restrictions reduce direct leakage of government dollars abroad, they cause government spending shocks to behave like negative labor supply shocks to the private sector, dampening the GDP response.

Keywords: Buy America, domestic content restrictions, government contracts, fiscal multiplier

*Acosta: miguel.acosta@wisc.edu Cox: lydia.cox@wisc.edu. We thank Jonathan Dingel, Aaron Flaaen, Ethan Ilzetzki, Amit Khandelwal, Adrien Matray, Emi Nakamura, Ricardo Reyes-Heroles, Peter Schott, Chenzi Xu, and audiences at the Federal Reserve Board, the Atlanta Fed Workshop on International Economics, the Chicago Workshop on International Economics, the Society for Economic Dynamics Annual Meeting, the Stanford Institute for Theoretical Economics 2024 Workshop on Trade and Finance, the Federal Reserve Bank of Dallas Conference on International Economics, Depaul University, Duke University, the HeiTüHo Workshop on International Financial Markets, the University of Tübingen, the University of British Columbia, and the NYU Trade and Spatial Conference.

1 Introduction

The U.S. federal government is the largest single buyer of goods and services in the world, spending an average of $600 billion per year. For almost a century, the government has wielded this power by placing domestic content restrictions on its own spending through a policy called the "Buy American Act" (BAA). Policymakers claim that by ensuring that "when the federal government spends taxpayer dollars, they are spent on American made goods by American workers and with American-made component parts," federal spending will bolster U.S. manufacturing, increase the resilience of U.S. supply chains, strengthen the industrial base, and ultimately create a "better," stronger economy (Biden Administration, 2021a,b). By providing a source of stable demand for domestic products, policymakers envision that firms may shift production to occur within U.S. borders. Despite having been in place for almost a century, there is little to no empirical evidence of the macroeconomic effects of such policies.

In this paper, we ask whether Buy American restrictions lead to stronger macroeconomic effects of government spending. Making progress on this question is complicated by the fact that Buy American restrictions are not directly observable in even the most detailed data on U.S. government spending. We propose a methodology based on a comprehensive understanding of Buy American regulations to identify government spending that is more- and less-constrained by these regulations. We validate our methodology in a novel firm-level dataset. We then apply our methodology to federal spending at the state level. We show that less-constrained contract dollars have a higher cross-sectional fiscal multiplier than more-constrained contract dollars. This finding may seem surprising, since Buy American constraints keep government spending inside the country. However, Buy American constraints are just that: constraints on production that, on net, outweigh the benefits of spending locally on total domestic production. This leads to a lower multiplier relative to unconstrained spending. We provide empirical evidence for this mechanism, then formalize the intuition in the context of a model of U.S. regions that accommodates our empirical setting.

Our first contribution is one of measurement: We leverage the complexity of the Buy American legislation—loopholes, amendments, and exemptions—to identify government spending that is more and less subject to Buy American restrictions. We apply our methodology to federal defense procurement contracts between 1979 and 2000, and the universe of procurement contracts between 2001 and 2019. These data provide detailed information on the contracting relationships between U.S. firms and the government, but provide almost no window into how much firms rely on imported inputs for their production and, therefore, how much they will be affected by Buy American restrictions. However, the general Buy American restrictions are not applied uniformly to all government contracts. For example, smaller contracts are exempt from Buy American restrictions altogether, while certain military purchases are regulated more stringently than non-military purchases. We propose three ways for grouping contracts into "constrained" and "unconstrained" contracts[1], where each of the three approaches is based on a different aspect of the Buy American legislation. We take advantage of details of the legislation in order to ensure that the contracts in these groups are similar to each other, up to the extent to which they are subject to Buy American restrictions. Therefore, we argue that we have identified variation in the strength of domestic content restrictions that is plausibly exogenous to firm-level and macroeconomic variables, allowing us to identify the causal effects of the Buy American restrictions.

We validate our measures of constrained and unconstrained spending by pairing the procurement data with data on firm-level import shipments from S&P Panjiva. These shipments are recorded on bills of lading that firms file with U.S. Customs on their maritime imports. Thanks to the high-frequency nature of the bill-of-lading data, we can track the real-time import behavior of firms when they receive government contracts that are more- and less constrained by the Buy American policy. We show that a firm's import response after receiving a constrained contract is significantly smaller than its import response to an unconstrained contract *for the same product*. These results provide evidence that the "constrained vs. unconstrained" breakdowns of contracts do indeed cause different import behavior by firms. We take this as evidence that, despite being imperfect measures of Buy American restrictions, our breakdowns still identify meaningful differences in the strength with which these restrictions are applied.

Having shown that we can effectively measure variation in the Buy American policy in the government contracts data, we turn to our primary question: How do Buy American restrictions influence the macroeconomic effects of government purchases? To address this question, we follow a longstanding literature that measures the "fiscal multiplier"—the percentage increase in gross domestic product (GDP) that results when government spending is increased by one percent of GDP. In particular, we rely primarily on a cross-sectional fiscal multiplier approach in which we exploit variation in regional military procurement associated with changes in aggregate defense spending. In our specification, we distinguish between the effects of government dollars that are more constrained by Buy American restrictions, and the effect of government dollars that are less constrained. We find no significant relative multiplier for constrained spending, and a relative multiplier of just under 2 for unconstrained

[1]What we refer to as "unconstrained" contracts going forward are more precisely referred to as "less constrained" relative to the constrained comparison group.

spending. We confirm that these effects are not driven by inherent differences in the industrial composition of constrained versus unconstrained contracts, nor by inter-state spillovers that might challenge our identifying assumptions.

On one hand, our finding that government dollars that are constrained by Buy American restrictions have a lower relative multiplier than unconstrained dollars can be rationalized by the fact that Buy American restrictions impose a constraint on production which, almost by definition, implies a lower level of output per dollar. On the other hand, the fiscal multiplier is about gross *domestic* product, and constrained dollars are forced to remain in the United States, while unconstrained dollars can "leak" out through imports, potentially putting downward pressure on GDP. We develop a framework that can rationalize our empirical results. Specifically, we model a monetary union, which allows us to calculate cross-sectional fiscal multipliers. In the model, firms produce both for the private sector and the government. The government imposes domestic content restrictions on a portion of its spending. Consistent with our empirical findings, we show that a shock to constrained government spending generates a lower relative multiplier. Key to this finding is that constrained spending shocks act like labor supply shocks with respect to the rest of the economy: As labor is reallocated to meet the government's demand, wages rise. This lowers output in the private sector and induces private-sector firms to demand more imports. In other words, though it is indirect, the constrained shock also induces "leakage." We provide empirical evidence consistent with this mechanism: The responses of wages, employment, and inflation are all higher for constrained spending than unconstrained spending.

This paper contributes to several strands of literature. First, there has been theoretical analysis on domestic content restrictions in trade agreements, from the seminal paper of Grossman (1981) to the more recent work of Grossman et al. (2023), who study whether it is optimal for governments to subsidize particular sources of inputs to mitigate supply chain risk. Additionally, rules of origin requirements in trade agreements have been studied both theoretically and empirically. Conconi et al. (2018), for example, show that the rules of origin requirements in NAFTA led to a reduction in imports of intermediate goods from third countries. Head et al. (2023) point, as we do, to the potentially counterproductive nature of rules of origin requirements. Other recent work by Allcott et al. (2024) studies the effects Buy American restrictions in the electric vehicle industry. Though related in spirit, we ask a fundamentally different question in this paper: whether domestic content restrictions on government purchases affect the fiscal transmission mechanism.

The body of work studying domestic content restrictions on government procurement, specifically, is smaller, particularly for the United States. In a policy brief, Hufbauer and Schott (2009) argues that Buy American provisions in the American Recovery and Reinvestment Act would do more damage to American employment than it would help. Dixon et al. (2018) use a computable general equilibrium model to evaluate the employment effects of the Buy American Act, and find that BAA is ineffective in generating employment. In a recent working paper by Bombardini et al. (2024), the authors use a quantitative trade model to assess the costs of Buy American provisions, and conclude that Buy American rules, and in particular the recent tightening of the Buy American constraints, will create few jobs at a relatively high cost. Our approach differs from these papers: by identifying Buy American restrictions directly in the data, we are able to produce causal estimates of the effects of the Buy American Act without relying on the assumptions of a particular structural model. There is also work by, for example, Herz and Varela-Irimia (2020) and García-Santana and Santamaría (2021, 2024) who study the effects of home bias in government contracting in Europe.

On the fiscal policy side, several papers have used different subsets of the government contract data to study fiscal transmission—e.g., Fisher and Peters (2010), Dupor and Guerrero (2017), Demyanyk et al. (2019), Auerbach et al. (2020), Cox et al. (2024)—as well as different aspects of the federal procurement process and its effects on firms—e.g., Hebous and Zimmermann (2021), Bajari et al. (2014), Ryan (2020), Warren (2014), Kang and Miller (2022), Bandiera et al. (2009), Gagnepain et al. (2013), and di Giovanni et al. (2022). Relative to these papers, we study a different fiscal transmission mechanism—how domestic content restrictions affect fiscal transmission in the United States.

Methodologically, our paper builds on the cross-sectional fiscal multiplier approach reviewed by Chodorow-Reich (2019), and in particular the identification strategy of Nakamura and Steinsson (2014). Our findings suggest that the regulations applied to government spending play a large role in determining the effect of that spending on output. In terms of data, we have benefited greatly from the work of Flaaen et al. (2021) in understanding the strengths and weaknesses of the Panjiva data that we rely on in our analysis. To our knowledge, ours is the first paper attempting to isolate variation in the import content of procurement contracts, and then measure how regulations on that content affects firms and the macroeconomy.

The rest of the paper proceeds as follows: In Section 2, we provide background and institutional details on the Buy American Act, in Section 3, we describe our primary data sources, and how we use them to measure variation in Buy American Act stringency across contracts; in Section 4 we present exercises that validate our measures, which we then use in our estimates of the cross-sectional fiscal multiplier in Section 5. In Section 6 we develop a framework that can rationalize our empirical result, before concluding in Section 7.

2 Background on the Buy American Act

The United States has had domestic content restrictions on federal procurement in place for almost a century. The flagship policy—The Buy American Act—was signed on President Herbert Hoover's last day in office in 1933. The policy, which has changed very little since its inception, has two basic requirements: when procuring materials, supplies, or manufactured end-use products, the government must purchase products that are manufactured

1. "substantially all" from materials that are produced in the United States, and
2. in the United States.

Determining whether a product qualifies as being domestically produced in accordance with Buy American standards involves some nuanced interpretation. If the government directly purchases unmanufactured products, those products must be mined or produced in the United States to qualify.[2] Instead, manufactured end-use products satisfy the basic requirements if they are manufactured in the United States and the cost of all intermediate components that are mined, produced, or manufactured in the United States exceeds 65 percent. (This threshold was increased by the Biden administration from 55 percent, and is scheduled to tighten further, to 75 percent by 2029.)

On top of these general requirements, there exist both exceptions and additional regulations for some procurement spending. For example, Buy American Act requirements only apply to procurement contracts above a certain monetary threshold—the "micro-purchase threshold," which is currently set at $10,000. There are also exceptions if awarding the contract to a domestic manufacturer would come at an "unreasonable cost"[3]. On the flip-side, there are certain cases in which the BAA requirements are even more stringent. The Berry Amendment, which has existed since the 1940s, provides that certain goods procured by the Department of Defense (DOD) must be entirely produced in the United States (i.e., 100 percent of the cost of intermediate components must be produced or manufactured domestically). These are just a few examples of the many complicated details of the Buy American Act legislation.

[2]There are additional regulations for products consisting wholly or predominantly of iron or steel. Specifically, the cost of domestic iron or steel must be greater than 95 percent of the cost of all components to qualify. See 48 C.F.R. §25.003.

[3]To determine whether the cost is "unreasonable," when a domestic bid for a contract is not the lowest bid, the procuring agency must add a cost penalty to the foreign bid (inclusive of duties) before determining which offer is the lowest. The size of the penalty is currently between 20 and 50 percent, depending on the firm and contracting agency. Currently, there is a 20 percent price penalty when the lowest domestic offer is from a large business, a 30 percent penalty when the lowest domestic offer is from a small business, and a 50 percent penalty when the contracting agency is part of the Department of Defense (DOD). After the penalty is applied, if the foreign offer is still lower than the domestic offer, the agency may award the contract to the foreign bidder at the initially proposed price.

Until recently, there had been very few changes to the Buy American Act since its inception. During the Great Recession, some Buy American requirements were built into the American Recovery and Reinvestment Act. The first Trump administration also used executive orders to reaffirm aspects of the Buy American Act that apply to iron and steel production. The largest change to the legislation in 70 years, however, occurred during the Biden administration. In May 2022, the Biden administration announced that the domestic content threshold would be gradually increased from its then rate of 55 percent to 75 percent by 2029. By ensuring that "when the federal government spends taxpayer dollars, they are spent on American made goods by American workers and with American component parts," the stated goals of the policy are to bolster U.S. manufacturing, increase the resilience of supply chains and create a stronger industrial base, and simply create a "better," stronger economy. As explained on the government "Made in America" website[4], the expectation is that "strengthening Made in America policies will send clear market signals to give suppliers confidence that manufacturing in the U.S., with America's workers, will provide greater opportunities. Ensuring Made in America laws are implemented clearly and consistently across government will support domestic suppliers." The goal of this paper is to provide empirical evidence as to whether Buy American restriction, indeed, lead to the stronger macroeconomic effects of government spending that policymakers promote.

3 Measurement of Buy American Restrictions

Against the background of the Buy American Act laid out in the previous section, we now describe how we are able to identify government purchases that are more and less subject to Buy American restrictions since, as far as we know, there is no direct measurement of this aspect of federal spending.

3.1 Government Procurement Contracts

Our measures of government spending are based on federal government procurement contracts. When the federal government purchases goods and services, it goes through a multi-stage solicitation and bidding process that ends when the contract is awarded to one of the bidding firms. In the contracts data, we see detailed information about every contract action that is taken—from government purchases of military equipment to office supplies to administrative services. From 1979–2000, we have access to contracts issued by the military

[4]This website—`www.madeinamerica.gov`—was launched during the Biden Administration in conjunction with a new "Made In America Office."

with values over \$25,000 (\$10,000 before 1984), collected on forms DD-350 and available from the National Archives.[5] Starting in 2000, we have access to the universe of government spending contracts, which are available from USAspending.gov and are described in detail by Cox et al. (2024). When analyzing these data over a sample that covers both the pre- and post-2000 data, we subset the latter data to only military spending, to make it consistent with the former.[6]

For each contract transaction, we observe a fairly comprehensive set of characteristics. Most relevant for our study, we see the date the transaction occurred (the `action_date`), the awarding agency (e.g., the Department of Defense or the Department of Homeland Security), the recipient firm, the dollar amount obligated from the government to the recipient firm, the six-digit NAICS code of the product or service being purchased (in later data), an additional product code known as the Federal Supply Classification (FSC), and the primary location in which the contract is to be fulfilled (at the zip code level). The granularity of the contracts data allows us to flexibly aggregate spending into government purchases of different types of goods or spending in different firms, industries, and locations. One example from previous work is that of Boehm (2016), who aggregates the contracts to the industry-year level and whose procedure we use to infer SIC industries from the "product service" or "federal supply" codes that accompany each contract before 1989. We rely on data aggregated to either the firm $\times$ sector $\times$ month level (in Section 4) or state $\times$ year level (in Section 5).

3.2 Direct Measurement of Buy American Restrictions

Assessing the extent to which the government consumes imported content and, in turn, is impacted by Buy American regulations, is challenging, even with the most disaggregated government procurement data. The government can purchase imported goods either directly, by contracting with a foreign supplier, or indirectly, by contracting with a U.S. manufacturer that uses imports as intermediates in production. The federal procurement data provide only limited information on these two channels. In the post-2000 data from USASpending.gov, for each contract, we observe the country in which the procured good or service originated.[7] If the government imports goods directly, it should be indicated by this variable. Table 1 shows that direct imports by the government from foreign countries are virtually non-existent:

[5]Available at `https://www.archives.gov/research/electronic-records/reference-report/federal-contracts`.

[6]Specifically, we follow Nakamura and Steinsson (2014) and restrict our sample to contracts awarded by the Department of the Army, Department of the Navy, Department of the Air Force, and the Defense Logistics Agency. These defense contracts represent roughly half of all contracts by count and around 60 percent by value.

[7]We can gather similar information from the pre-2000 sample in a less direct manner, so for simplicity we restrict our summary statistics to the post-2000 data here.

Table 1: Federal Purchases: Country of Origin

Country Code	Share of Contracts (%)
USA	90.13
IRQ	1.42
AFG	1.36
KWT	0.92
DEU	0.74
KOR	0.58
JPN	0.58
ARE	0.41
GBR	0.41
CAN	0.34

NOTE. This table shows the top ten countries, from 2001-2021, across which contract dollars are distributed.

Over the period from 2001-2021, 90 percent of contracted dollars are spent on goods or services produced in the United States, roughly 5 percent of contracted dollars are spent on goods or services produced in Iraq, Afghanistan, and Kuwait—spending on items used by deployed forces in those countries—and the remaining 5 percent is spread out over a number of countries. Ultimately, directly contracting with foreign companies does not appear to be a pervasive feature of government procurement.

The second way that the government can consume imported good is if it purchases products that are manufactured in the U.S. using imported intermediate inputs. In the contracts data, there is a variable called `place_of_manufacture`, which provides some information to this effect. There are eleven possible designations for this variable, describing whether the good was manufactured inside or outside of the United States, and if outside, what exception to the BAA rules permitted that purchase to happen. Consistent with Table 1, around 90 percent of contracted dollars are for items manufactured in the United States. Roughly 5 percent are manufactured outside of the U.S. for use outside of the U.S. (e.g., spending on products for deployed forces in the Middle Eastern countries shown in Table 1), and the remaining 5 percent are manufactured outside the U.S. for various reasons. The full breakdown is shown in Table 2.

At face value, the data suggest that the direct import content of federal procurement is very low. On indirect imports, the binary nature of what is indicated in the data masks an understanding of how much government purchases actually rely on intermediate inputs. An anecdotal example sheds light on the problem. A company called "Capps Shoe Company, LLC," that produces uniform dress shoes for the military and other law enforcement agencies,

Table 2: Federal Purchases: Place of Manufacture

Code	Place of Manufacture	Share of Contracts (%)
D	Mfg. in U.S.	89.74
E	Mfg Outside U.S. for use Outside U.S.	4.53
B	Mfg Outside U.S.	3.39
L	Mfg Outside U.S.—Qualifying Country (DOD)	0.94
A	Mfg in U.S., More than 50 perc. Foreign Content	0.84
J	Mfg Outside U.S.—Domestic Nonavailability	0.26
G	Mfg Outside U.S.—Trade Agreements	0.15
H	Mfg Outside U.S.—Commercial IT	0.1
F	Mfg Outside U.S.—Retail	0.03
I	Mfg Outside U.S.—Public Interest	0.01
K	Mfg Outside U.S.—Unreasonable Cost	0.01

Note. This table shows the percent of contracts that are manufactured in the U.S., or manufactured outside of the U.S. for one of many designated reasons. We exclude contracts in which the place of manufacture is missing, and we exclude contracts for goods that are "not manufactured end products" (where BAA restrictions do not apply).

has been a consistent supplier of the federal government since 2002, receiving contracts from the Department of Defense, the Department of Homeland Security (which houses the U.S. Coast Guard), and the Department of Transportation. Specifically, the text descriptions of the contracts awarded to Capps reveal that these government agencies purchase *Men's Footwear*, *Women's Footwear*, and *Women's Pumps*. Consistent with the contract descriptions, Capps website reveals that the company produces three styles of shoes: the "Angel Series" (a women's pump), the "Capital Series" (a unisex laced dress shoe), and the "Colonel Series" (similar to the Capital Series). More importantly, the website provides additional information about each style of shoe—relevant screenshots from their website are shown in Figure 1. The Angel Series pumps are described as "Made in the USA with imported parts", the Capital Series are described as "Made in the USA", and the Colonel Series are described as "Made in the USA and Berry Amendment Compliant." As we will discuss below, the Berry Amendment is an amendment to the Buy American Act that makes the domestic content restrictions even stricter for certain purchases. Despite these three separate labels, which clearly indicate differing levels of imported intermediates, the `place_of_manufacture` variable for all contracts awarded to Capps Footwear, LLC is reported as "Manufactured in the U.S." As a result, we believe there is more variation in the imported input content of government purchases than the contracts data suggest. In the rest of this section, we describe our efforts to capture some of the variation in the government's imported intermediate usage through an indirect approach.

Figure 1: Product Descriptions of three Capps Shoes Styles

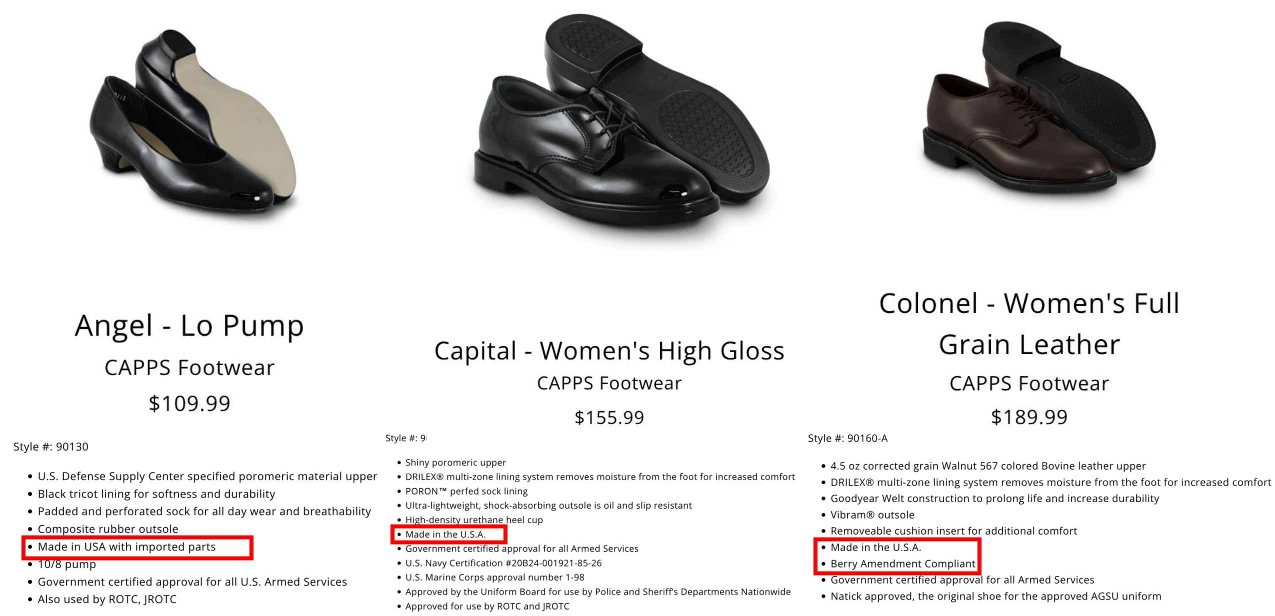

NOTE. This figure shows screenshots of product descriptions for three styles of Capps shoes (`https://originalfootwear.com/collections/capps-footwear`, last accessed October 2022).

3.3 Indirect Measurement of Buy American Constraints

In this subsection, we describe an alternative approach that allows us to capture variation in Buy American Act restrictions in order to measure the effects of those restrictions on firms and the macroeconomy. Specifically, we exploit complexities in the law that generate quasi-exogenous variation in the stringency of BAA restrictions. Leveraging these details, we can distinguish between contract dollars that are more and less constrained by domestic content requirements, allowing us to test the impact of these constraints on firms and macroeconomic variables. In particular, there are three exceptions and amendments to the Buy American Act that we leverage: the micro-purchase threshold, the Berry Amendment, and the WTO Government Procurement Agreement. We describe each in detail below.

The Micro-Purchase Threshold The Federal Acquisition Regulations describe a micro-purchase as "an acquisition of supplies or services, the aggregate amount of which does not exceed the micro-purchase threshold."[8] The micro-purchase thresholds for supplies and professional services are currently $10,000. Importantly, the Buy American Act only applies to purchases valued at more than the micro-purchase threshold. This means that purchases

[8] See FAR 2.101.

of, say, $10,001 must be produced with domestic content. Purchases of only $9,999, however, are free to be imported or to be produced in the United States using an unrestricted level of imported inputs. A potential concern is that contracting firms purposefully bid under the Micro-Purchase Threshold to evade Buy American requirements, but in Appendix A.1, we show that there appears to be little evidence of this kind of bunching.

The Berry Amendment The Berry Amendment is a statutory requirement that requires purchases of certain goods by the Department of Defense be made of 100 percent domestic content. The Amendment was first passed as part of the 1941 Fifth Supplemental Department of Defense Appropriations Act in order to protect the domestic industrial base in the time of war. The Amendment was included in subsequent defense appropriations acts until it was made permanent in 1994 and eventually codified into law in 2002.[9] The Berry Amendment applies to a certain set of products (classified by Federal Supply Classification codes), including textiles, clothing, footwear, food, hand or measuring tools, stainless steel flatware, and dinnerware. When these covered products are sold to the Department of Defense (or a Defense sub-agency), they must be made with 100 percent domestic content. When the same products are sold to a non-Defense agency, however, they must only satisfy the standard Buy American regulations.

The WTO Government Procurement Agreement The World Trade Organization's (WTO's) Agreement on Government Procurement (GPA) is a plurilateral agreement within the framework of the WTO. Currently 22 parties are official members of the Agreement, and several other countries participate either as observers or have initiated accession negotiations. The goal of the GPA is to mutually open government procurement markets among its members. Participating countries agree on certain procurement activities that are covered by the agreement, and for these activities, domestic content restrictions are not permitted. The United States' coverage schedule has both a product-based component and a threshold-based component. Only certain products are covered by the GPA (again, designated by FSC codes), and among covered products, only those above a given value threshold are covered. The threshold changes every few years, but typically ranges between $150,000 and $200,000. Products that are *not* covered by the agreement (even if above the threshold) include the products covered by the Berry Amendment (described above) and many types of transportation equipment, weapons, and specialty metal products. We collect the specific products and threshold values from coverage schedules that are published by the WTO.[10]

[9] See 10 U.S.C. 2533a.

[10] These schedules can be found at `https://www.wto.org/english/tratop_e/gproc_e/gp_app_agree_e.htm`.

Table 3: Definitions of Constrained and Unconstrained Contracts

Amendment or Exception	Constrained Definition	Unconstrained Definition
Micro-Purchase Threshold	Procurement contract value $\in [\text{MPT}_t, \text{MPT}_t + \$2500)$, where MPT_t is the micro-purchase threshold in year t	Procurement contract value $\in [\text{MPT}_t - 2500, \text{MPT}_t)$, where MPT_t is the micro-purchase threshold in year t
Berry Amendment	Products (FSC codes) covered by the Berry Amendment, purchased by the Department of Defense	Products (FSC codes) covered by the Berry Amendment, purchased by a non-Defense agency
WTO GPA	Products (FSC codes) *not* covered by the GPA OR contracts valued below GPA_t, where GPA_t is the relevant value threshold in year t	Products (FSC codes) covered by the GPA and contract valued above GPA_t, where GPA_t is the relevant value threshold in year t

For our purposes, when a contract is both a relevant product and is above the value threshold, contracting firms are permitted to use imports from any of the 22 member countries freely. When a contract is not covered by the Agreement, however, the usual Buy American restrictions apply.

Constrained and Unconstrained Contracts Using the three policy amendments and exceptions described above, we create three pairs of spending series, each with a "constrained" component and an "unconstrained" component. (We will refer to the "unconstrained" series, though in most cases these contracts are not completely unconstrained, just less constrained than the constrained series.) The definitions, described in Table 3, are based either on the product or service being purchased (designated by FSC code) or the contract value. For the Micro-Purchase Threshold, we define "constrained" contracts to be those just above the micro-purchase threshold (between the threshold and the threshold plus $2500).[11] "Unconstrained" micro-purchase contracts are those valued just below the threshold, (between the threshold and the threshold minus $2500). For the Berry Amendment, we define constrained contracts as contracts for products (FSC codes) covered by the Berry Amendment and purchased by the Department of Defense. Unconstrained contracts are defined as products that are covered by the Berry Amendment, but purchased by a Non-Defense agency, so the Berry Amendment does not apply. Lastly, for the WTO Government Procurement Agreement, we define unconstrained contracts as those covered by the GPA—this includes products (FSC codes) that are covered by the agreement when the contract value is above the relevant GPA threshold. Constrained GPA contracts include both products that

[11] The micro-purchase threshold is time-varying, but ranges from $2500 - $10,000 in our sample.

are not covered by the GPA and products that are covered, but when contract values are less than the GPA threshold.

Our claim is that using any of these breakdowns, we are able to construct series of government spending in the data that vary primarily with respect to how stringent their Buy American restrictions are. In the next section, we provide evidence validating these measures, which we then use to test the macroeconomic effects of the Buy American Act.

4 Measure Validation Using Bill of Lading Data

In this section, we provide empirical evidence that the "constrained" and "unconstrained" spending series that we have constructed appear to measure spending that is more and less constrained by the Buy American Act. To do this, we merge data on firm-level import shipments with the procurement contract data. We show that firms appear to import less when they receive what we define as a "constrained" contract relative to when they receive an "unconstrained" contract for the same 6-digit NAICS industry.

4.1 Bills of Lading Data from S&P Panjiva

We combine the government contracts data with publicly-accessible data on firm-level imports, collected by the U.S. Customs and Border Patrol (CBP). While CBP does not host the data, other third parties process and make the bill of lading data available for a cost. We use data supplied by S&P Panjiva, and are sensitive to the effort of Flaaen et al. (2021) who describe the strengths and weaknesses of the data. On the strengths side, we take advantage of the fact that the data show shipments at the firm level, which we are able to match to the procurement data. Another strength of the data is that they are also available in high frequency—by way of analogy with financial markets, the data are available at a "tick" frequency. Here, a tick is a shipment, though we aggregate these data at the firm $\times$ month level. The data are available starting in 2007.

One weakness of the Panjiva data is that they are an unreliable source of import *values*, though volumes are measured well. While Panjiva makes an effort to impute the products (with Harmonized System (HS) codes) contained in a shipment, to then assign those products to unit values, we have concluded anecdotally (as Flaaen et al. suggest) that this imputation contains a fair amount of measurement error. In many cases, the HS codes are imputed using the text description of imported products. In one case, S&P classified a shipment by a company that produces *vacuum measuring instruments* as a *vacuum cleaner*. Table A.1 in Appendix A.2 shows imports by HS code associated with government contracts in different

NAICS sectors. The first column, for example, shows which HS codes are imported when a firm receives a contract for "aircraft manufacturing." Clearly, while some of the products make sense, others (like "toys, including model planes") are examples of where the text-based imputation goes wrong. Therefore, we are unable to reliably impute a *dollar amount* of imports that are caused by a government contract, and largely avoid doing so in our analysis, relying on import volumes instead.

4.2 Measure Validation

We start by merging the universe of federal procurement contracts with the firm-level data on import shipments from Panjiva using S&P's "Business Entity Cross Reference Service Dataset." For each government contract transaction, we make use of the date the contract was agreed upon (or, the `action_date`), the recipient firm, the value of the contract, the six-digit NAICS sector that the contract is for, and the Federal Supply Classification. Using these characteristics for each firm × NAICS 6 industry × month, we aggregate the set of contracts that the firm receives that are constrained (according to one of our three definitions) and the set that are unconstrained. We then track the real-time import behavior of firms when they receive these constrained and unconstrained government contracts. We aggregate these to the firm-month level.

Using our merged dataset, we then estimate the following:

$$m_{f,t+h} - m_{f,t-1} = \delta_{f,i,h} + \delta_{t,h} + \beta_h^c g_{f,i,t}^c + \beta_h^u g_{f,i,t}^u + \varepsilon_{f,i,t,h}, \tag{1}$$

where $m_{f,t}$ is the inverse-hyperbolic sine (IHS) of import weight (in kg) (from Panjiva) by firm f in month t. Our primary independent variables of interest are $g_{f,i,t}^c$—the IHS of deflated constrained contracts by firm f in industry i, month t, and $g_{f,i,t}^u$—the IHS of deflated unconstrained contracts by firm f in industry i, month t.[12] We include firm × industry and time fixed effects. We only include firms whose lifetime contracts are majority manufacturing and i is a manufacturing industry. Our sample period is 2007–2022.

Of particular importance is the fact that we are comparing constrained and unconstrained contracts to the *same firm* for a product in the *same industry*.[13] Holding these characteristics

[12] We deflate these series by the U.S. personal consumption expenditure price index. We use the IHS because there are many zeros in our dataset. Technically, the estimates here are not elasticities because the IHS is not invariant to scale. However, in Appendix A.3, we discuss alternative specifications that lead us to conclude that the magnitudes presented here can be interpreted as elasticities.

[13] While the contract data are available at the firm-industry level, the import data are not. Thus, we have repeated values of the left-hand side variable in our baseline specification. While our clustered standard errors account for these repeated values, we have also estimated a specification in which we collapse the contract data to the firm level. The estimates are similar to our baseline, and are in Appendix A.4.

Figure 2: Elasticity of Imports to Constrained vs Unconstrained Contracts

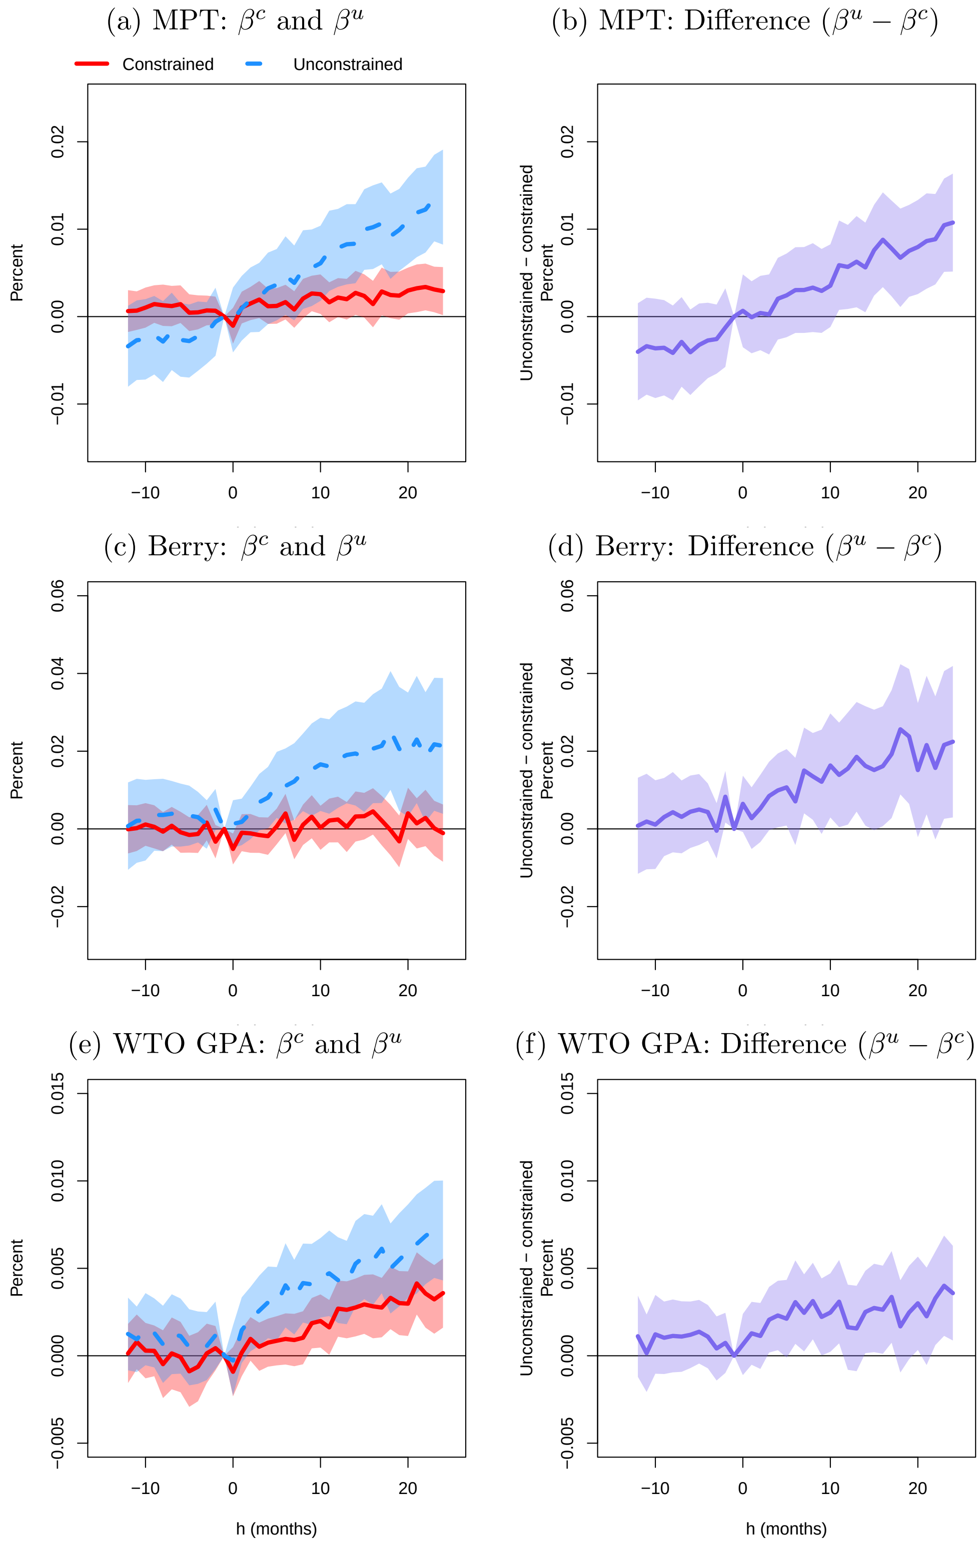

NOTE. The left panels show estimates of β^c_h (constrained) and β^u_h (unconstrained) from equation 1, for $h \in \{-12, \dots, 24\}$. The right panels show estimates of the differences. The shaded regions represent 95% confidence intervals constructed using standard errors that are clustered by firm-industry and month.

fixed and making use of the careful definitions of constrained and unconstrained contracts presented above, our claim is that the differential response of imports to these contracts is plausibly exogenous with respect to other fundamentals of the contract. In other words, while the level of the estimated β coefficients may potentially be biased if, for example, more productive firms import more and receive more contracts, our claim is that whether or not they receive a *constrained* contract in a month is independent of their productivity. This allows us to interpret the difference between β^c and β^u as the causal effect of the policy. We estimate this regression for all three definitions of "constrained" and "unconstrained" spending. Due to the limitations of the Panjiva data described earlier, we note that these coefficients are estimated elasticities of firm import *volumes* to a constrained and unconstrained contract *dollars*. These are informative about elasticities of import *values* to the extent that product composition and import prices remain unchanged.

The results are plotted in Figure 2. In the left column of the figure, we plot the individual elasticities of imports to constrained and unconstrained contracts, β^c and β^u, respectively. In the right column, we plot the difference between the two elasticities. In both figures, the shaded regions indicate 95 percent confidence intervals. The top panel of the figure shows the results when we use the micro-purchase threshold (MPT) to define constrained and unconstrained contracts. Upon receipt of a constrained contract for a particular product (NAICS 6 industry), we see a very limited response of imports in the Panjiva data—indicated by the relatively flat, solid red line. Receipt of an unconstrained contract by the same firm and for the same product, however, generates a positive and statistically significant import response in the Panjiva data over the subsequent two years.[14] And most importantly, the panel on the right shows that the response of imports to constrained contracts is statistically significantly lower than the import response to unconstrained contracts. The middle and bottom panels show results from estimating the same specification, but using the Berry Amendment definition of constrained and unconstrained (middle panel) and WTO government procurement agreement definition (bottom panel).

The results are similar for all three formulations. Of particular note are the Berry Amendment results in the middle panel. Recall, the Berry amendment requires that when certain products are purchased by the Department of Defense, they are required to be made of 100 percent domestic content. When those same products are purchased by a non-Defense

[14] Though our focus is not on the magnitude of these elasticities, a back-of-the-envelope calculation suggests that they are in a reasonable range: In our 2007–2022 sample, the portion of contracts we consider averaged about 250 billion dollars annually. Splitting this in half, we have about 125 billion dollars in constrained contracts, and 125 billion dollars in unconstrained contracts. Over this period, total business sales (FRED mnemonic `TOTBUSSMNSA`) averaged about 16 trillion dollars annually. So, contracts of each type constitute a 0.008 share of sales, consistent with the elasticities of 0.005 to 0.02 that we estimate.

agency, only the regular Buy American requirements apply. In the regression, we are comparing the import response of a firm, f, when it receives a contract for a Berry product, i, from the Defense Department (constrained) to the response of the same firm when it receives a contract for the same Berry product from a non-Defense agency (unconstrained). As predicted, the constrained Berry contract elicits zero import response while we see a slight positive response for the unconstrained contract. In what follows, we explore whether these constraints have implications for the macroeconomic effects of government spending.

5 BAA and the Fiscal Multiplier

Using our procedure for identifying constrained and unconstrained government spending, we turn to our primary question: How do Buy American restrictions influence the macroeconomic effects of government purchases? To address this question, we follow the literature that measures the fiscal multiplier. In particular, we estimate, in the words of Chodorow-Reich (2019), a "geographic cross-sectional fiscal multiplier"—the effect of an increase in government spending on output in one region of the United States relative to another. To rule out potential reverse causality, we build on the instrumental variables approach of Nakamura and Steinsson (2014), who exploit variation in state military procurement associated with aggregate changes in aggregate defense spending. Relative to typical cross-sectional fiscal-multiplier regressions, we distinguish between the effects of government dollars that are constrained by Buy American restrictions and the effect of government dollars that are unconstrained, or less constrained.

5.1 Baseline Estimates

Our baseline empirical specification is a cross-sectional fiscal multiplier regression:

$$\frac{Y_{s,t} - Y_{s,t-2}}{Y_{s,t-2}} = \beta^c \left(\frac{G^c_{s,t} - G^c_{s,t-2}}{Y_{s,t-2}} \right) + \beta^u \left(\frac{G^u_{s,t} - G^u_{s,t-2}}{Y_{s,t-2}} \right) + \Gamma_1 \mathbf{X}_{s,t} + \alpha_s + \gamma_t + \varepsilon_{s,t}. \quad (2)$$

The left-hand side is the two-year percent change in real per-capita GDP in state s in year t.[15] The right-hand side includes the two-year percent change in defense spending in state s in year t. We split defense spending into two components: "constrained" spending, captured

[15] Our data on state-level GDP come from the Bureau of Economic Analysis (BEA). Post-1997 data are available through their API—the nominal GDP series code is `SAGDP2N`. Until 1997, the data are available for download from `https://apps.bea.gov/regional/zip/SAGDP_SIC.zip`—in that file, the nominal GDP series code is `SAGDP2S`. The data from the two sources are constructed slightly differently, so we rescale the later data so that the values in 1997 are identical to the earlier data.

by $G^c_{s,t}$ and "unconstrained" spending, captured by $G^u_{s,t}$. State and time fixed effects absorb any state-specific trends in output and government spending, and aggregate shocks. We control for two lags (from $t-3$ to $t-2$ and from $t-4$ to $t-3$) of both government-spending variables and the outcome variables in $\mathbf{X}_{s,t}$, following the suggestion of Ramey (2021). We deflate GDP and government spending by the U.S. personal consumption expenditure price index, and put them in per-capita terms using state-level population data from the Census Bureau. The sample period for our baseline regression starts in 1996 and ends in 2019, and we consider only contracts that are designated as "manufactured end-use products"—the set of contracts to which the Buy American Act applies.[16] We cluster standard errors by state and year, as recommended by Majerovitz and Sastry (2023).

One concern with equation 2 is that, given the political nature of government spending, the federal government may award contracts to certain states in times of economic need. To address this concern, we instrument for the percent changes in $G^c_{s,t}$ and $G^u_{s,t}$ with percent changes in aggregate constrained or unconstrained spending from $t-2$ to $t-1$, each interacted with an indicator for state s.[17] The identification of β^c and β^u thus comes from the fact that states have different sensitivities to changes in aggregate (constrained or unconstrained) defense spending. To interpret these estimates as causal, we assume that changes in aggregate defense spending are exogenous to state economic conditions. In other words, we assume that the government does not engage in military build ups in order to help particular economies, and that the allocation of spending between more- and less-constrained contracts during military build-ups reflects the needs of the military, and not the needs of particular states. We also include in $\mathbf{X}_{s,t}$ two lags of our instruments—the growth of national spending in each category interacted with a state fixed-effect.

We allocate each defense contract based on how constrained that contract is by domestic content restrictions. What we call "unconstrained" purchases, here, are those that were covered by the WTO's Government Procurement Agreement. "Constrained" purchases are any remaining defense procurement purchases for manufactured end-use-products. This breakdown allows us to effectively split aggregate contract spending in half. The other breakdowns that we described in Section 3—those governed by the Berry Amendment and Micro-purchase threshold—provide us series of contracts that represent only a small portion of government procurement—too small to have measurable macroeconomic effects. The

[16] Specifically, in the post-2006 data, there is a variable that indicates whether a contract is "not a manufactured end use product." We select a set of product or service (PSC) codes that are *never* labeled as "not a manufactured end use product," and only keep contracts that are designated under those PSC codes.

[17] The timing of our specification is: the shock occurs between $t-2$ and $t-1$, and we estimate a two-year "integral multiplier" (Ramey and Zubairy, 2018) over $t-2$ to t. We discuss a fully dynamic specification below.

Table 4: State Shares of Constrained and Unconstrained Contracts

Constrained	Share	Unconstrained	Share
Maine	0.82	Utah	0.77
Connecticut	0.77	New Mexico	0.77
Missouri	0.69	North Dakota	0.78
Vermont	0.68	District of Columbia	0.78
Iowa	0.68	Wyoming	0.79
Mississippi	0.67	Montana	0.79
New Hampshire	0.66	Hawaii	0.82
Minnesota	0.65	Nebraska	0.83
Washington	0.6	Wisconsin	0.83
Arkansas	0.6	Oklahoma	0.87

NOTE. The table above shows the states with the highest shares of contracts for manufacturing end-use-products coming from constrained contracts (left) and unconstrained contracts (right). The shares are computed over our baseline sample period from 1996 to 2019.

modern-day version of the WTO GPA went into force in 1996, which is why we choose that year as the beginning of the sample in our baseline regressions. To provide a sense of which states we expect to be more sensitive to the different contract types, in Table 4, we list the states for which constrained and unconstrained contracts represent the highest shares of total contracts, based on our definition.

Table 5 contains our baseline estimates of equation (2). The first row shows the effect of increasing government spending by one percent of GDP when that spending is relatively constrained by domestic content restrictions. The second row shows the effect when the spending is relatively unconstrained. We find a cross-sectional multiplier of -0.86 for constrained spending, which is not precisely estimated, and 1.88 for unconstrained spending, which is statistically significant. To put these estimates into perspective, the corresponding estimates of Nakamura and Steinsson for spending that is not disaggregated by Buy American Act constraints is 1.43. Our estimate is based on a different set of contracts and years, uses different standard errors, and includes controls for lags so, in Appendix B.1, we confirm that our data and sample period yields a similar multiplier of 1.48 when not disaggregated. We also test the null hypothesis that the coefficient on constrained contracts is larger than the coefficient on unconstrained contracts, and can reject the null with a p-value of 0.01, suggesting there is a statistically significant difference in the GDP response to the different types of spending.

Table 5: Fiscal Multipliers by Constrained and Unconstrained Spending

	GDP
Constrained ($G^c_{s,t}$)	-0.86
	(0.78)
Unconstrained ($G^u_{s,t}$)	1.88
	(0.72)
Observations	1129
p(uncon<con)	0.01
Sample	1996–2019

NOTE. This table presents IV estimates of equation (2) for state real per-capita GDP. We test the null hypothesis that the coefficient on constrained is larger than the coefficient on unconstrained spending, and report the associated p value in the row labeled "p(uncon<con)." Standard errors clustered by state and year are in parentheses.

5.2 Threats to Identification

Industry Composition. Our interpretation of the results presented in Section 5.1 is that domestic content restrictions dampen the macroeconomic effects of government spending. That interpretation depends on whether domestic content restrictions are the main distinguishing feature between what we call "constrained" and "unconstrained" contracts. The construction of our baseline breakdown, based on the WTO GPA, opens the possibility that the set of constrained contracts has a different industry composition than the set of unconstrained contracts. This is because contracts covered by the GPA—the unconstrained contracts—are contracts that (1) exceed a certain threshold and (2) are for the purchase of a specific set of products, Ω. This means that the set of unconstrained contracts excludes contracts for any goods that are not in Ω, which may skew the industry composition of unconstrained contracts relative to constrained contracts.

Reassuringly, there is substantial overlap in the industries making up both the constrained and unconstrained series in our baseline. This is reflected in Table 6, which shows the ten industries receiving the largest amounts of constrained and unconstrained contracts over the course of our sample. Industries that receive a large portion of overall U.S. federal procurement—such as Aircraft Manufacturing—show up in both the constrained and unconstrained part of the sample. There are still slight differences in the industry composition of our two series, however, so we present two additional exercises to help mitigate the concern that composition is driving our results.

First, we consider an alternative breakdown of contracts that holds the set of products fixed. The alternative series of constrained spending includes contracts for products in Ω

Table 6: Top Industry Recipients of Constrained and Unconstrained Contracts

Constrained Contracts

NAICS Code	Description
336411	Aircraft Mfg
336611	Ship Building, Repairs
336414	Guided Missile and Space Vehicle Mfg
336413	Oth Aircraft Parts and Equip. Mfg
334511	Navigation, Guidance, and Oth. Instrument Mfg
336412	Aircraft Engine, Engine Parts Mfg
541330	Engineering Services
336992	Military Armored Vehicle, Tank, Parts Mfg
332993	Ammunition
334220	Radio, TV Broadcasting, Wireless Comm. Equip. Mfg
Unconstrained Contracts	
541330	Engineering Services
336411	Aircraft Mfg
541715	Research in Physical, Engineering, and Life Sciences
324110	Petroleum Refineries
236220	Commercial and Institutional Bldg. Construction
336413	Oth Aircraft Parts and Equip. Mfg
561210	Facilities Support Services
336611	Ship Building, Repairs
336992	Military Armored Vehicle, Tank, Parts Mfg
488190	Support Activities for Air Transportation

with contract values that are between zero dollars and the threshold. The alternative series of unconstrained spending includes contracts for products in Ω with contract values that are between the threshold and double the threshold. Despite being a much smaller set of contracts, we detect macroeconomic effects that are comparable and consistent with our baseline estimates. In Appendix B.2, we describe this exercise in more detail.

Second, we make use of the fact that the WTO Government Procurement Agreement has changed over time. The agreement was originally established in 1979 as the "Tokyo Round Government Procurement Code," but was renegotiated and expanded as part of the Uruguay Round of the General Agreement on Tariffs and Trade, the predecessor of the WTO. The more modern version was signed in 1994 and implemented in 1996. We perform a placebo test of sorts in which we classify contracts by the *modern* GPA definition of thresholds and covered products, in the pre-1996 data. The idea is that before 1996, those classifications did not directly demarcate more- and less-constrained contracts. We then estimate our baseline regression, but separately estimate coefficients for constrained and unconstrained spending pre- and post-1996. The results are shown in Appendix B.3. In the pre-1996 data, we cannot statistically distinguish β^c from β^u. Post-1996, we continue to reject the null that $\beta_c > \beta_u$. This suggests that it is the policy change, not the industry composition of the spending series, driving our baseline results.

Inter-state spillovers. Inter-state spillovers present a challenge to the identification of cross-sectional fiscal multipliers in general, and are potentially more of a concern in our setting. Consider the experiment that identifies cross-sectional multipliers: A national military build-up begins and, because of existing characteristics (like Connecticut's industrial composition), more contract dollars (as a percent of GDP) flow to Connecticut than New York. We then compare how GDP changes in Connecticut relative to New York. If, however, Connecticut contractors use inputs from New York to fulfill their contracts, the relative increase in Connecticut's GDP will be understated. In our setting, we consider the effect of spending a *constrained* dollar in Connecticut. Because the contracting firm in Connecticut is required to fulfill that contract using domestic inputs, it may be even more likely to source inputs from New York. This would lead to an even larger understatement of the regional fiscal multiplier for constrained contracts, which could explain the smaller multiplier that we estimate for constrained spending relative to unconstrained spending.

We have several reasons to believe that inter-state spillovers are not driving our primary results. First, there is evidence from previous work which suggests that inter-state spillovers from procurement contracts are limited. Nakamura and Steinsson (2014), for example, show that the value of shipments from defense-oriented industries to the government within states moves one for one with military procurement in that state. In other words, a dollar sent from the government to Connecticut tends to be fulfilled in value by shipments from Connecticut back to the Government. Auerbach et al. (2020) show that fiscal multipliers dissipate quickly with distance: a contract given to a firm in one city has a high local multiplier, but a multiplier of 0 in locations that are more than 50 miles away.

Second, we provide further supporting evidence using our data that inter-state spillovers are unlikely to be driving our results. Specifically, we aggregate our data from the state-level to the Census-region level. A Census region is a group of neighboring states. Again, the concern is that inter-state spillovers may be leading us to overstate the difference between β^c and β^u. Under the assumption that inter-region spillovers are less likely than inter-state spillovers, if inter-state spillovers were driving the difference, we should expect a smaller difference using regional data. In terms of the example, the contract dollar given to Connecticut but spent in New York would be part of the same Census region, so the role of interstate transfers should be diminished in the regional data. Instead, as we show in Appendix B.4, the difference between β^c and β^U actually grows when we aggregate to the regional level. Taken together, the evidence suggests inter-state spillovers are likely not driving our results.

Buy-American at the national level. The identification in our baseline IV estimation of Equation (2) depends on the assumption that the allocation of national spending between more- and less-constrained contracts during military build-ups reflects the needs of the military, and not the needs of particular states. An alternative approach with a weaker assumption is to instrument for changes in $G^c_{s,t}$ and $G^u_{s,t}$ with state-indicators interacted with national military procurement that is not split by the extent of domestic content restrictions. We present those estimates in Appendix B.1. The results using this alternative instrument are similar to our baseline.

Many instruments. Our IV estimation of Equation (2) includes two instruments for every state: we interact a state indicator with constrained spending at the aggregate level, and unconstrained spending at the aggregate level. The presence of many instruments can complicate inference. We therefore present an alternative approach based on a Bartik instrument in Appendix B.5.[18] For each state, the instrument for constrained spending is constructed by multiplying the change in aggregate spending by the average fraction of constrained spending to GDP in that state between 1979 (the start of our data) and 1995 (the year before the implementation of the GPA). While the standard errors on each of the point estimates are much larger, we continue to find statistical evidence that favors the hypothesis that constrained spending leads to a smaller cross-sectional multiplier than unconstrained spending.

Dynamic Responses Our baseline specification is akin to a cumulative two-year cross-sectional "integral" fiscal multiplier, to use the language of Ramey and Zubairy (2018). In Figure B.1 in the appendix, we present estimates of a version of Equation (2) in which we replace the left-hand side with $(Y_{s,t+h} - Y_{s,t-2})/Y_{s,t-2}$ for $h = 0, \ldots, 5$. These integral fiscal multipliers suggest that our baseline estimates are fairly stable out to five years beyond the initial stimulus. The estimated multipliers become somewhat larger for unconstrained spending, though statistically it would be challenging to detect any meaningful difference.

[18] We also estimated our Bartik-instrument specification using the "re-centering" method of Borusyak and Hull (2023)—the results were nearly identical to those shown in Appendix B.5.

5.3 Additional Outcomes

Our main focus in this paper is on the role that Buy-American restrictions play in determining the fiscal multiplier. In Section 6, we develop a model that can rationalize our empirical findings. To help guide the modeling exercise and to provide a more-complete view of the empirical effects, we present estimates of the effects of government spending on additional economic outcomes in Table 7.

The first two columns in Table 7 focus on the labor market. The first column shows the effects on the employment-to-population ratio, and the second shows the effects on real compensation per worker.[19] While the coefficients are statistically indistinguishable from zero, there is a statistically significant difference in the effect that constrained and unconstrained spending have on these variables. Qualitatively, the effects are in the opposing direction from the effects on GDP. Constrained spending leads to a larger effect on both employment and average compensation than unconstrained spending.

The final column shows the effect on inflation.[20] Here, again, the individual coefficients are not indistinguishable from zero, and the evidence for differential effects is weaker. However, a similar qualitative finding emerges: Constrained spending leads to more inflation than unconstrained spending.

Taken together, the response of the labor market and inflation are more consistent with how one would expect spending that is constrained to remain "in the country" to operate. What can explain the difference between these responses and that of GDP? Intuitively, the Buy American Act requires that each unit of output be produced with a higher share of domestic inputs—including domestic labor. Given that the BAA restrictions appear to be binding (based on the evidence presented in Section 4), it is unsurprising that prices also rise more in response to constrained spending. In Section 6, we develop a model to shed light on the mechanism driving the GDP response to the Buy American constraints that is consistent with the labor market and price effects. The key insight from the model is that, while BAA spending leads to higher wages and employment, there is a spillover of higher wages through the labor market that increases costs for all firms and, on net, decreases production and raises prices. We formalize this intuition next.

[19]In the EPOP specification, the left-hand side is $\frac{E_{s,t}-E_{s,t-2}}{P_{t-2}}$, where $E_{s,t}$ is employment and $P_{s,t-2}$ is population in state s, year $t-2$. We take state-level employment from the Bureau of Labor Statistic's Quarterly Census of Employment and Wages. We define real compensation, $w_{s,t}$, as the ratio of real worker compensation divided by employment. Real worker compensation is from the same source as GDP (the BEA, described in footnote 15): the ratio of `SAGDP4` (compensation of employees) to `SAGDP8` (GDP deflator). The left-hand side in that specification is $\frac{w_{s,t}-w_{s,t-2}}{w_{s,t-2}}$.

[20]Our data on state-level inflation are from Hazell et al. (2022). Their data do not cover all states and end earlier than our baseline sample period. From their inflation data, we create a price index for each state $P_{s,t}$, then include the two-year percent change $\left(\frac{P_{s,t}-P_{s,t-2}}{P_{s,t-2}}\right)$ on the left-hand side.

Table 7: Effects of Government Spending: Additional Outcomes

	EPOP	Comp. per worker	Inflation
Constrained ($G^c_{s,t}$)	0.15	0.55	0.33
	(0.09)	(0.48)	(0.19)
Unconstrained ($G^u_{s,t}$)	0.08	-1.72	-0.14
	(0.12)	(1.08)	(0.32)
Observations	1129	1129	714
p(uncon<con)	0.01	0.04	0.07
Sample	1996–2019	1996–2019	1996–2017

Note. This table presents IV estimates of equation (2) for additional outcome variables, described at the top of each column. The additional variables are described in the text. All estimates are responses to government spending shocks that are one percent of GDP. The response of employment is in percentage points, and the other responses are in percent. We test the null hypothesis that the coefficient on constrained is larger than the coefficient on unconstrained spending, and report the associated p value in the row labeled "p(uncon<con)." Standard errors clustered by state and year are in parentheses.

6 Potential Mechanisms

We showed in the previous section that government dollars that are constrained by Buy American restrictions have a lower relative multiplier than unconstrained dollars. This could be explained by the fact that Buy American restrictions impose constraints on production. However, constrained dollars are required to remain in the country, which would tend to boost domestic production.

We now develop a framework that can rationalize our empirical results. We model a monetary union in the spirit of Galí and Monacelli (2008), which allows us to calculate cross-sectional fiscal multipliers. Relative to typical monetary union models, we extend the supply side of the economy to accommodate our empirical setting. Specifically, in the model, firms produce for both consumers and the government. The government imposes a domestic content restriction on a portion of its spending. Consistent with our empirical findings, we show that a shock to constrained government spending leads to a lower cross-sectional multiplier. Key to this finding is that constrained government spending shocks act like labor supply shocks to consumption-good production: As labor is reallocated to meet the government's demand, wages rise. This lowers output in the consumption-good sector, and also induces consumption-good firms to demand more imports. The net effect on GDP reflects the joint impact of the negative spillovers to consumption-good production outweighing the positive effects of restricting imports directly.

6.1 A Macro Model of the Buy American Act

Our model has two regions, a and b that belong to a monetary and fiscal union. When comparing regional outcomes, region a will be the region in which the government spending shock occurs and region b will represent the rest of the economy. In each region there are two types of production: production that satisfies Buy American restrictions (the "BAA sector"), which we will denote with tildes, and production that does not satisfy Buy American restrictions (the "consumption-good sector"). The Buy American goods are purchased only by the government, while the standard consumption goods are purchased both by the government and by households. We describe the key features of the model briefly here, and provide further details in Appendix C.

6.1.1 Households

Region a has a continuum of household types indexed by x, where x indicates the type of labor supplied by the household. Households maximize lifetime utility over consumption and leisure:

$$\mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t \left[\frac{C_t^{1-\frac{1}{\sigma}}}{1-\frac{1}{\sigma}} - \chi \frac{\mathcal{L}_t(x)^{1+\frac{1}{\nu}}}{1+\frac{1}{\nu}} \right].$$

Households consume a CES composite of the goods produced in regions a and b, given by:

$$C_t = \left[\phi_a^{\frac{1}{\eta}} C_{at}^{\frac{\eta-1}{\eta}} + \phi_b^{\frac{1}{\eta}} C_{bt}^{\frac{\eta-1}{\eta}} \right]^{\frac{\eta}{\eta-1}}.$$

The goods in each region, C_{at} and C_{bt}, are also CES composites of varieties produced in each region, which are indexed by z, with elasticity of substitution θ. Each household allocates labor between production in two sectors: Buy American production (where labor is $\widetilde{L}_t(x)$) and consumption-good production ($L_t(x)$) such that

$$\mathcal{L}_t(x) = \widetilde{L}_t(x) + L_t(x).$$

Households are indifferent between working in the two sectors: Wages are the same in each, and it is costless to transition from one sector to the other. Households in region b have an analogous problem. A fraction n of the national population live in region a, and $1-n$ live in region b. Households in the two regions purchase one-period bonds that ensure perfect risk-sharing between the regions.

6.1.2 Government

The government purchases two types of goods from each region: those produced according to the consumption-good technology (G_{at} and G_{bt}), and goods produced using a technology that adheres to Buy American requirements ($\widetilde{G}_{at}, \widetilde{G}_{bt}$). Each of the four types of government spending follow AR(1) processes, and we assume that Buy-American spending shocks are uncorrelated with consumption-good spending shocks. The government's demand for each firm follows the same form as the household's, with CES preferences over varieties z and elasticity of substitution θ. The central bank sets the nominal interest rate according to an inertial Taylor (1993) rule that includes national CPI inflation and GDP.

6.1.3 Firms

There is a continuum of firms in each region, indexed by z. Firm z specializes in the production of differentiated good z to produce output $y_{at}(z)$. The production function of each firm is given by:

$$y_{at}(z) = \phi \left[\alpha L_t(z)^\rho + (1-\alpha) M_t(z)^\rho\right]^{\frac{1}{\rho}} \tag{3}$$

where $M_t(z)$ is an imported intermediate, $L_t(z)$ domestic labor, and $\frac{1}{1-\rho}$ is the elasticity of substitution between inputs. We abstract from modeling the market for imports directly, and assume that inputs are available in infinite supply at the exogenous price μ_t from the rest of the world. We assume that firms are subject to nominal rigidities as modeled by Calvo (1983) and Yun (1996), and can thus only change their price each period with probability κ. Following Woodford (2003), we assume each firm is part of an industry x with many other firms that use the same type of labor and set prices at the same time. We assume labor is mobile across sectors, but immobile across regions, and firms take the industry wage, $w_t(x)$ as given.

Each firm also produces a good that is compliant with Buy-American restrictions. The production technology for these goods is identical, except that the government imposes a fixed tariff τ on imports that is set to ensure that the import share of material costs is below the Buy American restriction. There is no revenue collected from this tariff, it is simply a wedge such that from the perspective of domestic producers selling the Buy American good, the price of imports goes from μ_t to $\tau\mu_t$.

6.2 Solution and Calibration

We log-linearize the model's equilibrium conditions around a zero-inflation steady state. We solve the model using first-order perturbation methods. The complete set of equilibrium conditions are in Appendix C.

Many of the structural parameters in our model are standard in the New-Keynesian literature: The model is calibrated to a quarterly frequency with the discount factor $\beta = 0.99$; the intertemporal elasticity of substitution $\sigma = 2$; and the probability that a firm can reset its price in any quarter is $(1 - \kappa) = 0.25$. We set the coefficients on the monetary policy rule to be consistent with the IV estimates in Table 5 of Coibion and Gorodnichenko (2012), with $(\rho_r, \rho_y, \rho_\pi) = (0.8, 1, 3)$. We follow Nakamura and Steinsson (2014) in setting inter-regional parameters: the elasticity of substitution between varieties $\theta = 7$, the elasticity of substitution between regions $\eta = 2$, the home-bias parameters $\phi_a = 0.69$, and the size of the region receiving the shock $n = 0.1$. We set the Frisch elasticity of substitution to $\nu = 0.75$, as suggested by Chetty et al. (2011).

We calibrate the government-spending parameters using our data. We set the quarterly autocorrelation of all government spending shocks to 0.95, consistent with the annual first-order autocorrelation of national real per-capita government spending in our data of 0.87.[21] We assume that 50% of government spending is subject to BAA requirements, roughly in line with our WTO-GPA delineation of constrained and unconstrained contracts in our data. Consistent with the national accounts, government spending is 20% of GDP.

To calibrate the parameters of the production functions (ρ, α, and τ), we use information from several sources. We begin by setting the elasticity of substitution between L and M to 2 (i.e., $\rho = 0.5$). This is consistent with the Armington elasticity estimated by (Boehm et al., 2023).[22] This implies that imported and domestic intermediates are substitutes. In the model, the ratio of the cost of a unit of labor to the cost of a unit of imports for consumption-good production is $^{w}/_{\mu}$, and $^{w}/_{\tau\mu}$ for BAA production. Thus, we normalize the common component, w/μ, to unity, then calibrate τ to differentiate the production types (described below).

To calibrate α, we rely on data on firm input costs from Eldridge and Powers (2018). From chart 5 of their paper, we take the 2015 factor input costs shares of labor (46%), imported materials (30%), and domestic materials (24%).[23] We consider these shares to be indicative of unconstrained cost shares, and thus use them to calibrate α using the consumption-good

[21] We arrive at this estimate via an AR(1) with a linear trend in a 1981–2019 sample.

[22] The mapping between the Armington elasticity and the elasticity between L and M in our model is not exact. In our model, M represents foreign intermediates, and L represents domestic labor. The Armington elasticity is more informative about the elasticity between foreign intermediates and domestic intermediates, the latter of which is not in our model. However, we have been unable to find estimates of the elasticity between foreign intermediates and domestic labor. To the extent that labor and intermediates are complements, one would use a lower value of ρ. We explore how our results change with a lower value of ρ in the text.

[23] The authors also report costs of other inputs, but since those do not map as cleanly into our model/BAA requirements, we exclude them from our calculations.

Figure 3: Responses to Government Spending Shocks

(a) Output (b) Imports (c) Wages

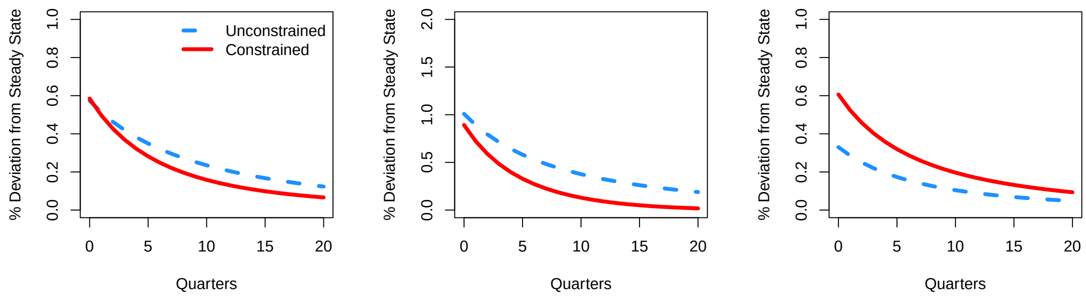

Note. This figure shows impulse responses to one-percent government-spending shocks in region a of the model. The blue lines correspond to shock to unconstrained (consumption-good) shocks, and the red lines correspond to constrained (BAA) shocks.

production function in the model. Specifically, one can show that the import share of costs in the model, which we will denote by φ,[24] is related to α by

$$\alpha = \frac{1}{1 + \frac{1}{\left(\frac{1-\varphi}{\varphi}\right)^{1-\rho}\left(\frac{w}{\mu}\right)^{\rho}}}. \tag{4}$$

The input cost data do not account for the fact that consumers can consume imports directly. To allow for this behavior in our calibration, we assume that final consumption is composed of 85% domestic goods, and 15% imported goods and adjust the input cost shares accordingly: The imported-materials share becomes $0.15 + 0.85 \times 0.30 = 0.405$, and the sum of the domestic input share becomes $0.85 \times (0.46 + 0.24) = 0.595$. Thus, with $\varphi = 0.405$ and other calibrated parameters, equation (4) implies $\alpha = 0.55$.

We calibrate τ based on the requirements of the BAA. Using the Eldridge and Powers (2018) input-cost shares, we reduce the share of imported materials and reallocate it to the domestic-materials share so that the share of imported materials in total material costs is 55%, consistent with the BAA requirement in place for most of our sample (or 75% when we re-calibrate to the more-stringent policy). This yields an imported-material share under the BAA (which we denote by $\widetilde{\varphi}$) of 0.24, a domestic material share of 0.30, and leaves the labor share unchanged (by construction) at 0.46. In the steady state of the model, τ can be expressed as

$$\tau = \frac{w}{\mu}\left(\frac{1-\widetilde{\varphi}}{\widetilde{\varphi}}\right)^{\frac{1-\rho}{\rho}}\left(\frac{1-\alpha}{\alpha}\right)^{\frac{1}{\rho}}.$$

So, given the values of calibrated parameters and $\widetilde{\varphi} = 0.24$, we have that $\tau = 2.2$.

[24] This is the ratio of μM (steady-state cost of imports) to SY (total costs: steady-state marginal costs per unit S, times steady-state output, Y.

Figure 4: Sectoral Decomposition of Impulse Responses

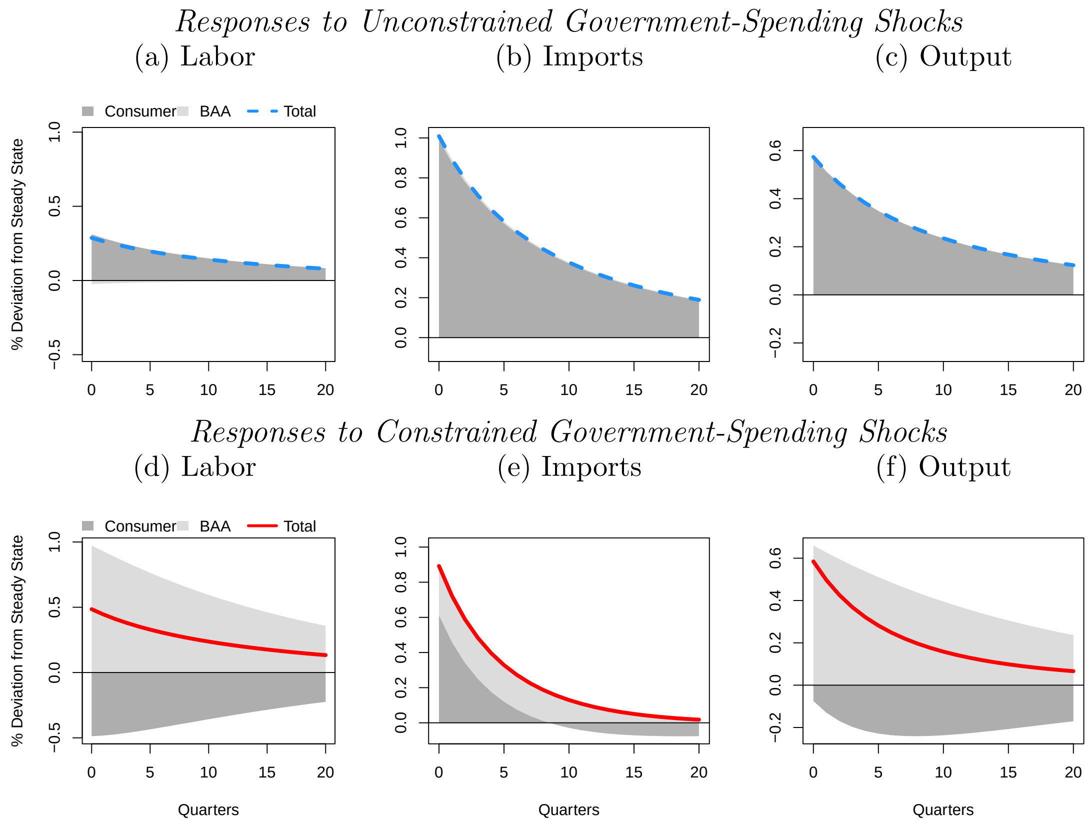

Note. This figure shows the response of total labor, imports, and output in region a in response to one-percent unconstrained (top) and BAA-constrained (bottom) government spending shocks in region a of the model. We define "total labor" (and, analogously, imports and output) as the sum of labor in the consumption-good and BAA sector (i.e., $\mathbf{L}_t \equiv L_t + \widetilde{L}_t$). We log-linearize the sum, so that $\widehat{\mathbf{L}}_t = \phi_L \widehat{L}_t + \widetilde{\phi}_L \widehat{\widetilde{L}}_t$, where ϕ_L and $\widetilde{\phi}_L$ denote steady-state shares of each type of labor, and hats denote percent deviations from steady state. The colored lines denote the response of $\widehat{\mathbf{L}}_t$, the dark-gray regions the response of $\phi_L \widehat{L}_t$, and the light-gray regions the response of $\widetilde{\phi}_L \widehat{\widetilde{L}}_t$,

6.3 Mechanism

To understand how the economy responds to government spending shocks, we plot impulse responses of output and imports in region a to unconstrained government spending (G_{at}) and constrained government spending ($\widetilde{G}_{at}$) in Figure 3. Two features stand out. First, while both shocks raise the respective form of government spending by one percent of output, the unconstrained spending translates into a higher general equilibrium response of output (Panel (a)). Second, imports (Panel (b)) respond by more to the unconstrained shock, but the response of imports to BAA spending is nearly as large.

Both of these features of the model can be explained by what happens in the labor market in region a. Panel (c) shows that wages increase by more in response to the constrained spending shock. This arises because the constrained spending shock requires more labor per unit of output. Because wages are the same across the two sectors, the increase affects the decisions for consumption-good production.

To illustrate this spillover, in Figure 4 we present a sectoral decomposition of the labor, import, and output responses to each type of government spending. The top panel shows that the response to the unconstrained shock are relatively straightforward: Both labor and imports increase in the consumption-good sector to fulfill the increase in government demand, and there is an accompanying increase in output. In this way, these responses behave like a standard demand shock. There are virtually no spillovers from the consumption-good sector into the BAA sector.[25]

In contrast, the lower panels of Figure 4 show that the responses to a constrained spending shock are more complicated. Panel (d) shows that there is a much larger response of labor supplied to the BAA sector relative to the unconstrained case. This increase is partially offset by a decrease in labor in the consumption-good sector, a consequence of the higher wages (Figure 3). These forces are reflected in the response of output to the constrained shock, which rises in the BAA sector, but is offset by a decline in the consumption-good sector (Panel (f)). The net impact is a smaller output response relative to the unconstrained shock, as shown in Figure 3. Unlike in the unconstrained case, the constrained shock manifests as a combination of a standard demand shock in the BAA sector and a contractionary labor supply shock in the consumption-good sector.

It is important to also consider the response of imports, which could offset the weaker response of output in estimating the response of GDP. Panel (e) shows considerable spillovers in the import response as well: Imports rise slightly in the BAA sector, but are accompanied by a large increase in imports in the consumption-good sector. Due to the increase in the price of labor, the constrained government spending shock induces an increase in imports in the rest of the economy. On net, the import response to the two types of shocks is similar—though it is indirect, the BAA shock also induces "leakage."

6.4 The Cross-Sectional Fiscal Multiplier

We now use the model to assess our empirical findings directly. To that end, we simulate a long time series of our model economy in the presence of government spending shocks of both types (constrained and unconstrained) in region a. We then estimate the cross-sectional multiplier: The response of region a relative to region b to each type of government

[25] Demand for the BAA good is exogenous, so by construction, output in the BAA sector does not changed. There is a small reallocation between imports $\widetilde{M}_t$ and labor $\widetilde{L}_t$ in the BAA sector as wages change.

Table 8: Regional and National Multipliers in the Model

Calibration	Cross-Sectional Multipliers Unconstrained	Cross-Sectional Multipliers Constrained	National Multipliers Unconstrained	National Multipliers Constrained
55% BAA Requirement				
Baseline	0.37	0.27	0.11	0.05
Low Frisch Elasticity	0.32	0.20	0.10	0.02
Low Intermediates Elasticity	0.40	0.13	0.11	−0.03
Neoclassical	0.33	0.21	0.23	0.05
75% BAA Requirement				
Baseline	0.37	0.24	0.12	0.02
Low Frisch Elasticity	0.33	0.15	0.10	−0.02
Low Intermediates Elasticity	0.41	0.01	0.11	−0.10
Neoclassical	0.33	0.16	0.24	−0.05

NOTE. This table shows estimates of the fiscal multiplier using simulated data from our model (we simulate the model for 100,000 quarters). In the columns labeled "cross-sectional multipliers," we estimate equation (2) via OLS (recall, government spending is exogenous in the model) in a scenario where only region a is receiving shocks to both types of government spending. In the columns labeled "national multipliers," we simulate shocks to each type of government spending that are perfectly correlated across the two regions, and report the response of aggregate GDP to each type of shock in a national analog of equation (2) (i.e., with no time fixed-effects).

spending. As in our baseline specification, the government shocks are scaled to be as a percent of output, and we estimate changes over two-year horizons in annual data.

Table 8 shows the results, in the column labeled "cross-sectional multipliers." The first row shows that the model qualitatively matches our main result: The cross-sectional multiplier for constrained spending is smaller than for unconstrained spending. In that table, we also explore various alternative calibrations to assess the role of different parameters. In the row labeled "Low Frisch elasticity," we decrease the Frisch elasticity from our baseline value of 0.75 to 0.5. This drives a wider wedge between the constrained and unconstrained multipliers. This is consistent with the mechanism described above—the lower labor supply elasticity implies that a constrained shock has to raise wages by more in order to attract the required workers. This exacerbates the spillover to private sector production.

The elasticity of substitution between the domestic input, L, and the imported input, M, is also an important determinant of the difference between the constrained and unconstrained fiscal multiplier. In the row labeled "Low Intermediates Elasticity," we lower the elasticity of substitution from $1/2$ to $1/3$. This again drives a bigger wedge between the two cross-sectional multipliers. The lower elasticity implies that, in response to government spending, wages rise by more since imports cannot offset the increase in demand for labor. This is worse for constrained spending, which requires a higher labor share per unit of output. Finally, our results hold qualitatively in the absence of sticky prices, as is displayed in the row labeled "Neoclassical."

In the bottom panel of the table, labeled "75% BAA requirement," we re-calibrate the parameter τ so that 75% of the cost of intermediates must be domestically sourced to be compliant with the BAA policy. This is consistent with the policy that will be in effect later this decade as a consequence of Biden-administration directives. While this has little consequence in our baseline calibration, it causes the difference between unconstrained and constrained spending to increase in the other calibrations we consider. In the case that domestic and foreign inputs become less substitutable, for example, the cross-sectional multiplier for constrained spending drops to zero.

6.5 The Aggregate Multiplier

In addition to shedding light on the mechanism driving our empirical findings, the model can be used to translate the cross-sectional fiscal multiplier that we estimate in the data and model into the more familiar aggregate multiplier. To compute aggregate multipliers in our model, we simulate government spending shocks that are perfectly correlated across the regions (making them "national shocks" as opposed to regional shocks), and estimate our fiscal multiplier on national data.

The results are shown in the column of Table 8 labeled "National Multiplier." In general, across the calibrations, constrained shocks tend to have smaller multipliers than unconstrained. In most cases the multiplier for constrained spending is essentially zero. Under the more-stringent 75% policy, the aggregate multipliers for constrained spending become negative in most of our calibrations. While the quantitative value of the aggregate multiplier is known to be sensitive to aggregate monetary policy, these results show that our empirical finding that unconstrained spending has a larger multiplier than constrained spending "survives" in the aggregate.

7 Conclusion

In this paper we study the macroeconomic effects of a longstanding policy that regulates the foreign content in government procurement: The Buy American Act. Our primary interest is in understanding whether Buy American Act restrictions lead to stronger macroeconomic effects of government spending, as policymakers hope. From an empirical standpoint, this is a challenge: Even in the most detailed data on U.S. government spending, Buy American regulations are hard to observe directly. We propose a new way of measuring Buy American Act restrictions that leverages the complexity of the law itself. The various loopholes, amendments, and exemptions allow us to identify variation in the strength of domestic content restrictions across contracts that is plausibly exogenous to firm-level and macroeconomic variables, allowing us to identify causal effects.

To address our primary question, we follow the longstanding fiscal multiplier literature. Exploiting variation in regional military procurement associated with changes in aggregate defense spending, we are able to show that government dollars that are less constrained by Buy American restrictions have a significantly larger cross-sectional multiplier than government dollars that are constrained by the policy. From a simple accounting perspective, our findings may appear surprising. However, we rationalize these findings with a general equilibrium model of our empirical setting. In the model, constrained spending shocks act like labor supply shocks with respect to the rest of the economy: By causing a larger reallocation of labor to meet the government's demand for a more labor-intensive good, private sector output falls and wages rise. Additionally, the increase in wages induces the private sector to demand more imports. In other words, though it is indirect, the constrained shock also induces "leakage." Additional cross-sectional evidence is consistent with this mechanism.

This paper investigates just one of many aspects of the Buy American Act—it's direct impact on GDP. Policymakers have stated other goals for this type of policy, such as increasing the resilience of supply chains. Further research is necessary to explore additional channels through which the Buy American Act can affect the macroeconomy.

References

Allcott, Hunt, Reigner Kane, Max Maydanchik, Joseph Shapiro, and Felix Tintelnot, "The Effects of "Buy American": Electric Vehicles and the Inflation Reduction Act," Working Paper 2024.

Auerbach, Alan, Yuriy Gorodnichenko, and Daniel Murphy, "Local fiscal multipliers and fiscal spillovers in the USA," *IMF Economic Review*, 2020, *68* (1), 195–229.

Bajari, Patrick, Stephanie Houghton, and Steven Tadelis, "Bidding for incomplete contracts: An empirical analysis of adaptation costs," *American Economic Review*, 2014, *104* (4), 1288–1319.

Bandiera, Oriana, Andrea Prat, and Tommaso Valletti, "Active and passive waste in government spending: evidence from a policy experiment," *American Economic Review*, 2009, *99* (4), 1278–1308.

Bartik, Timothy J, "Who benefits from state and local economic development policies?," 1991.

Biden Administration, "FACT SHEET: Biden-Harris Administration Issues Proposed Buy American Rule, Advancing the President's Commitment to Ensuring the Future of America Is Made in America by All of America's Workers," July 2021. Accessed: 2025-06-29.

_ , "President Biden to Sign Executive Order Strengthening Buy American Provisions, Ensuring Future of America Is Made in America by All of America's Workers," January 2021. Accessed: 2025-06-29.

Boehm, Christoph E., "Government Spending and Durable Goods," Working Paper 2016.

Boehm, Christoph E, Andrei A Levchenko, and Nitya Pandalai-Nayar, "The long and short (run) of trade elasticities," *American Economic Review*, 2023, *113* (4), 861–905.

Bombardini, Matilde, Andres Gonzalez-Lira, Bingjing Li, and Chiara Motta, "The Increasing Cost of Buying American," Technical Report, National Bureau of Economic Research 2024.

Borusyak, Kirill and Peter Hull, "Nonrandom exposure to exogenous shocks," *Econometrica*, 2023, *91* (6), 2155–2185.

Calvo, Guillermo A, "Staggered prices in a utility-maximizing framework," *Journal of Monetary Economics*, 1983, *12* (3), 383–398.

Chetty, Raj, Adam Guren, Day Manoli, and Andrea Weber, "Are micro and macro labor supply elasticities consistent? A review of evidence on the intensive and extensive margins," *American Economic Review*, 2011, *101* (3), 471–475.

Chodorow-Reich, Gabriel, "Geographic Cross-Sectional Fiscal Spending Multipliers: What Have We Learned?," *American Economic Journal: Economic Policy*, May 2019, *11* (2), 1–34.

Coibion, Olivier and Yuriy Gorodnichenko, "Why are target interest rate changes so persistent?," *American Economic Journal: Macroeconomics*, 2012, *4* (4), 126–162.

Conconi, Paola, Manuel García-Santana, Laura Puccio, and Roberto Venturini, "From final goods to inputs: the protectionist effect of rules of origin," *American Economic Review*, 2018, *108* (8), 2335–65.

Cox, Lydia, Gernot J Müller, Ernesto Pasten, Raphael Schoenle, and Michael Weber, "Big g," *Journal of Political Economy*, 2024, *132* (10), 000–000.

Demyanyk, Yuliya, Elena Loutskina, and Daniel Murphy, "Fiscal stimulus and consumer debt," *Review of Economics and Statistics*, 2019, *101* (4), 728–741.

di Giovanni, Julian, Manuel García-Santana, Priit Jeenas, Enrique Moral-Benito, and Josep Pijoan-Mas, "Buy Big or Buy Small? Procurement Policies, Firms' Financing, and the Macroeconomy," *Procurement Policies, Firms' Financing, and the Macroeconomy (February 1, 2022). FRB of New York Staff Report*, 2022, (1006).

Dixon, Peter B., Maureen T. Rimmer, and Robert G. Waschik, "Evaluating the effects of local content measures in a CGE model: Eliminating the U.S. Buy America(n) programs," *Economic Modelling*, 2018, *68*, 155–166.

Dupor, Bill and Rodrigo Guerrero, "Local and aggregate fiscal policy multipliers," *Journal of Monetary Economics*, 2017, *92*, 16–30.

Eldridge, Lucy P and Susan G Powers, "Imported inputs to US production and productivity: Two decades of evidence," Technical Report, Bureau of Labor Statistics 2018.

Fisher, Jonas DM and Ryan Peters, "Using stock returns to identify government spending shocks," *The Economic Journal*, 2010, *120* (544), 414–436.

Flaaen, Aaron, Flora Haberkorn, Logan T Lewis, Anderson Monken, Justin R Pierce, Rosemary Rhodes, and Madeleine Yi, "Bill of lading data in international trade research with an application to the COVID-19 pandemic," 2021.

Gagnepain, Philippe, Marc Ivaldi, and David Martimort, "The cost of contract renegotiation: Evidence from the local public sector," *American Economic Review*, 2013, *103* (6), 2352–83.

Galí, Jordi and Tommaso Monacelli, "Optimal monetary and fiscal policy in a currency union," *Journal of International Economics*, 2008, *76* (1), 116–132.

García-Santana, Manuel and Marta Santamaría, "Border Effects in Public procurement: the Aggregate Effects of Governments' Home Bias," 2021.

_ **and** _ , "Governments' Home Bias and Efficiency Losses: Evidence from National and Subnational Governments," Technical Report, CEPR Discussion Papers 2024.

Grossman, Gene M., "The Theory of Domestic Content Protection and Content Preference*," *The Quarterly Journal of Economics*, 11 1981, *96* (4), 583–603.

Grossman, Gene M, Elhanan Helpman, and Hugo Lhuillier, "Supply chain resilience: Should policy promote international diversification or reshoring?," *Journal of Political Economy*, 2023, *131* (12), 3462–3496.

Hazell, Jonathon, Juan Herre no, Emi Nakamura, and Jón Steinsson, "The slope of the Phillips Curve: evidence from US states," *The Quarterly Journal of Economics*, 2022, *137* (3), 1299–1344.

Head, Keith, Thierry Mayer, and Marc Melitz, "The unintended consequences of high regional content requirements," *Local Content Requirements*, 2023, p. 87.

Hebous, Shafik and Tom Zimmermann, "Can government demand stimulate private investment? Evidence from US federal procurement," *Journal of Monetary Economics*, 2021, *118*, 178–194.

Herz, Benedikt and Xosé-Luís Varela-Irimia, "Border effects in European public procurement," *Journal of Economic Geography*, 2020, *20* (6), 1359–1405.

Hufbauer, Gary Clyde and Jeffrey J Schott, *Buy American: Bad for jobs, worse for reputation*, Peter G. Peterson Institute for International Economics, 2009.

Kang, Karam and Robert A Miller, "Winning by Default: Why is There So Little Competition in Government Procurement?," *The Review of Economic Studies*, 2022, *89* (3), 1495–1556.

Majerovitz, Jeremy and Karthik Sastry, "How Much Should We Trust Regional-Exposure Designs?," Working Paper 2023.

Matray, Adrien, Karsten Müller, Chenzi Xu, and Poorya Kabir, "EXIM's exit: The real effects of trade financing by export credit agencies," Working Paper 2025.

Nakamura, Emi and Jon Steinsson, "Fiscal stimulus in a monetary union: Evidence from U.S. regions," *American Economic Review*, 2014, *104* (3), 753–92.

Ramey, Valerie A., "Comment," *NBER Macroeconomics Annual*, 2021, *35*, 232–241.

Ramey, Valerie A and Sarah Zubairy, "Government spending multipliers in good times and in bad: evidence from US historical data," *Journal of Political Economy*, 2018, *126* (2), 850–901.

Ryan, Nicholas, "Contract enforcement and productive efficiency: Evidence from the bidding and renegotiation of power contracts in India," *Econometrica*, 2020, *88* (2), 383–424.

Sims, Christopher A, "Solving linear rational expectations models," *Computational economics*, 2002, *20* (1), 1–20.

Taylor, John B, "Discretion versus policy rules in practice," in "Carnegie-Rochester conference series on public policy," Vol. 39 Elsevier 1993, pp. 195–214.

Warren, Patrick L, "Contracting officer workload, incomplete contracting, and contractual terms," *The RAND Journal of Economics*, 2014, *45* (2), 395–421.

Woodford, Michael, *Interest and Prices*, Princeton, NJ: Princeton University Press, 2003.

Yun, Tack, "Nominal price rigidity, money supply endogeneity, and business cycles," *Journal of Monetary Economics*, 1996, *37* (2), 345–370.

A Data and Measurement

A.1 Bunching around Micro-Purchase Threshold

A potential concern about using the Micro-Purchase Threshold for a regression discontinuity-type of analysis is that firms intentionally bid below the threshold to evade the Buy American restrictions. If this were the case, we would expect to see missing mass in the distribution of contract values just above the threshold. In Figure A.1, we show little evidence of this phenomenon. The left panel shows the distribution of contract values relative to the MPT when the MPT ranged from \$2500 to \$3500. If anything, the "bunching" goes in the wrong direction—there is an increase in contracts just *above* the threshold. In the right panel, we show the distribution of contract values relative to the MPT when the MPT was set at \$10,000 (it's current level). Here there is perhaps a small bit of missing mass to the right of the threshold, but it is minimal, and hard to distinguish from the choppiness of the rest of the distribution.

Figure A.1: Distribution of Contract Values Near MPT

(a) 2001-2016 (b) 2017-2022

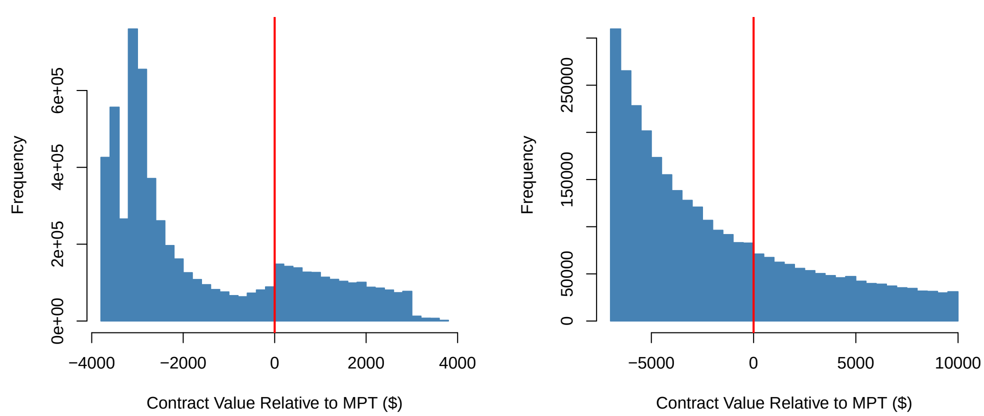

A.2 Imported Goods by Industry

In this appendix, we describe the relationship between the industry for which a firm receives a government contract, and the products that the firm imports. To that end, in Table A.1, we focus on the five 6-digit NAICS manufacturing or trade industries that received the most contracts since 2000. Unfortunately, we cannot see exactly which product a firm imports for which purpose, so instead we look at time-series correlations at the firm level in order to see which patterns emerge. Specifically, within each industry, we regress imports of each HS product by the firm on government contracts received by that firm in that industry. We call the resulting coefficient the sensitivity of imports of a given product j to contracts for industry i. We report the five HS products with the highest sensitivity that is statistically significant (at the 5% level or more) in the table.

Table A.1: Top Imported Goods by Industry of Government Contract

Aircraft manufacturing	Ship building/repair	Missiles and spacecraft	Navigation instruments	Aircraft parts
Toys, including model planes	Heat exchange units	Aircraft launching gear	X-ray tubes	Tapered roller bearings
Parts for electrical circuits	Parts of cooling equipment	Exercise equipment	Lamps	Cases/boxes/packages
Rubbers/plastic extruders	Cast articles of iron or steel	Bottles/flasks	Cases/boxes/packages	Parts for regulating instruments
Parts for gas filtering	Drive Axles	Other wooden furniture	Drive Axles	Parts for gas filtering
Chemical wood pulp	Motor starting equipment	Specialty motor vehicles	Toys, including model planes	Ovens/stoves/etc.

A.3 Variable Transformations in the Import-Elasticity Estimation

The main dependent and independent variables shown in Figure 2 are inverse-hyperbolic sine transformations (IHS) of the underlying variables.[26] We pursue that approach because our data on firm-level imports and, in particular, firm-industry level contracts contains many zero entries. One drawback of the IHS is that the resulting coefficients are not true elasticities, since the IHS is not invariant to scale. In this appendix, we discuss two alternative specifications to address these concerns. Both sets of results are shown in Figure A.2.

In panel I of Figure A.2, we show a simple alternative to our baseline transformation: we simply take the natural log of one plus the original variable. Let $M_{f,t}$ be import weight (in kilograms) for firm f in month t, and $G^i_{f,i,t}$ be contracts of type i (constrained or unconstrained) that firm f receives in month t for industry i. Then the left-hand side of equation (1) is $\ln(M_{f,t+h}+1) - \ln(M_{f,t-1}+1)$ and the independent variables are $\ln\left(G^i_{f,i,t}+1\right)$ for $i \in \{c, u\}$. The choice of adding 1 may be arbitrary, but it ensures that the zero observations contribute negligibly to the estimated effect. The estimates are nearly identical quantitatively to our baseline specification. Given these results, we feel comfortable interpreting the coefficients in Figure A.2 as elasticities.

In panel II, we explore an alternative transformation discussed by Matray et al. (2025) and references therein: The midpoint growth formula. We leave the government spending as in panel I, but replace the left-hand side with

$$\frac{m_{f,t+h} - m_{f,t-1}}{\frac{1}{2}\left(m_{f,t} + m_{f,t-1}\right)}.$$

This transformation allows imports of zero in either $t+h$ or $t-1$, and can be interpreted as a growth rate. We do not transform government spending, since the resulting number of observations is too small to allow us to estimate the effects with any precision.[27] While the estimated elasticities in panel II are attenuated by two thirds, and the difference between coefficients is estimated with less precision, we continue to find the same qualitative pattern across the three delineations of contracts: Constrained contracts are followed by fewer imports than unconstrained contracts.

[26] For a variable x, its inverse-hyperbolic sine is $\ln\left(x + \sqrt{x^2+1}\right)$.

[27] For reference, at $h = 0$, we have 13 million observations in our baseline specification (and in the log-log specification); 5 million in the midpoint-growth specification; and only 113,537 for WTO contracts, 60,263 for micro-purchase contracts and 1,230 for Berry-Amendment contracts in a specification with midpoint-growth for both the dependent and independent variables.

Figure A.2: Elasticity of Imports to Constrained vs Unconstrained Contracts

I. Log-Log Specification

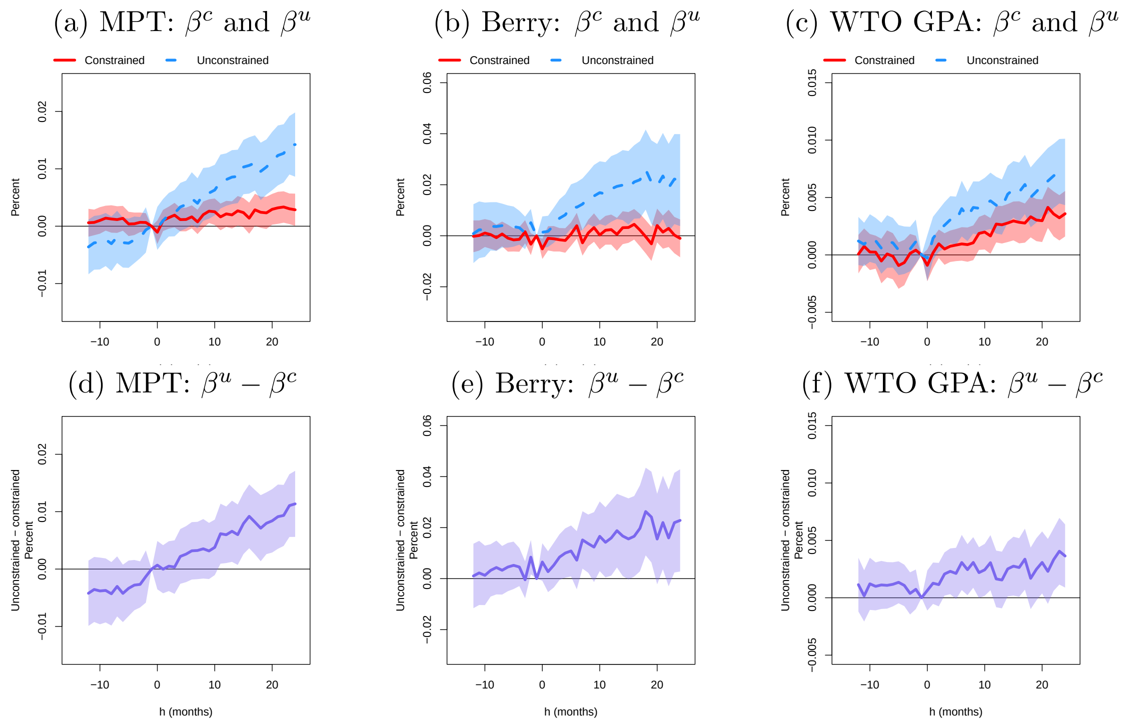

II. Midpoint-Growth Specification

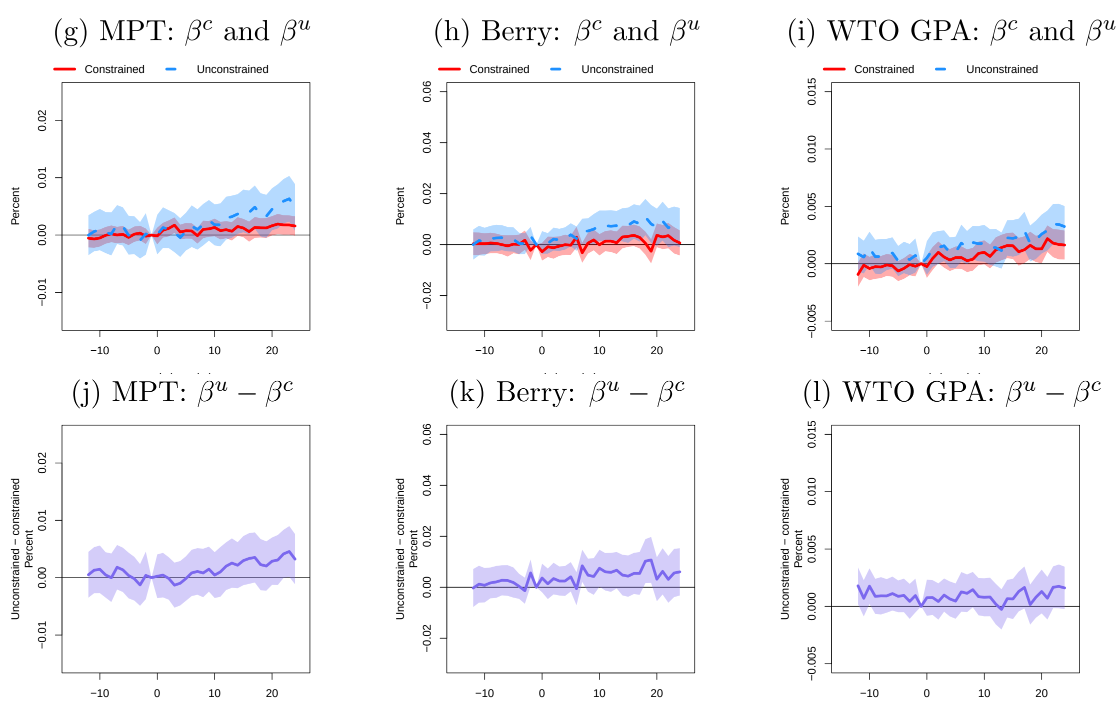

Note. This figure shows estimates of equation (1), for alternative transformations of the dependent and independent variables. See the discussion in Appendix A.3 for details. Shaded regions denote 95% confidence intervals using standard errors clustered by firm-industry and year.

A.4 Firm-Level Import-Elasticity Specification

Our main independent variables in equation (1) are at the firm-industry level, while our dependent variable is at the firm level. In this appendix, we estimate a similar specification, but at the firm level:

$$m_{f,t+h} - m_{f,t-1} = \delta_{f,h} + \delta_{t,h} + \beta_h^c g_{f,t}^c + \beta_h^u g_{f,t}^u + \varepsilon_{f,t,h}, \tag{A.1}$$

where $g_{f,t}^i$ is the inverse hyperbolic since of the sum of contracts of type i (constrained or unconstrained) received by firm f in month t, and $\delta_{f,h}$ is a firm fixed effect. Otherwise, everything is the same as in equation (1). The results are shown in Figure A.3, and are similar to our baseline.

Figure A.3: Import-Contract Elasticity with Firm-Level Contracts

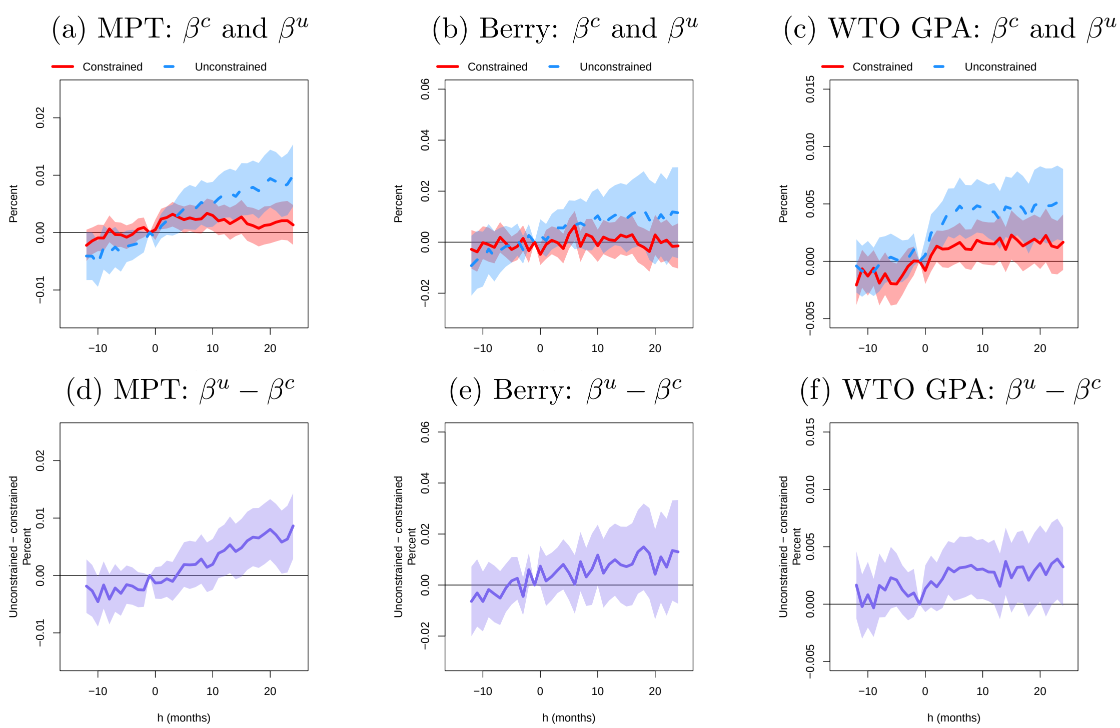

(a) MPT: β^c and β^u (b) Berry: β^c and β^u (c) WTO GPA: β^c and β^u

(d) MPT: $\beta^u - \beta^c$ (e) Berry: $\beta^u - \beta^c$ (f) WTO GPA: $\beta^u - \beta^c$

Note. This figure shows estimates of equation (A.1). Shaded regions denote 95% confidence intervals using standard errors clustered by firm and year.

B Fiscal Multiplier: Robustness

This appendix contains supporting tables for Section 5.

B.1 No disaggregation by type of spending

Our baseline specification is similar to that of Nakamura and Steinsson. Aside from the inclusion of two types of government spending, the specification differs in a couple of ways. First, our sample begins later (1996 instead of 1966). Second, we include lags and use two-way clustered standard errors, to address criticisms made of their specification over the years. In this appendix, we carry out a "replication" of their specification using our sample and econometric adjustments in order to assess the extent to which these are driving our results. In Table B.1, we estimate a specification analogous to that of Nakamura and Steinsson:

$$\frac{Y_{s,t} - Y_{s,t-2}}{Y_{s,t-2}} = \beta \left(\frac{G_{s,t} - G_{s,t-2}}{Y_{s,t-2}} \right) + \alpha_s + \gamma_t + \Gamma' \mathbf{X}_{s,t} + \varepsilon_{s,t}, \tag{B.1}$$

where $G_{s,t}$ is total defense procurement in state s in year t. We include in $\mathbf{X}_{s,t}$ two lags of $G_{s,t}$ and two lags of each instrument. The difference, relative to our baseline, is that we do not disaggregate government spending by how subject it is to Buy American restrictions. This serves as a benchmark for our disaggregated baseline estimates, and as a check that the different sample periods yield similar results. The point estimates for the maximal sample we can estimate using our data and specification (1983–2019) and for our baseline sample (1996–2019) are similar to the corresponding estimates of Nakamura and Steinsson, who estimate 1.4. The estimates in our baseline sample lie in between the estimates for constrained and unconstrained contracts presented in Table 5.

To address the concern that the allocation of national military spending into constrained and unconstrained contracts may be a response to state-specific factors, we consider alternative instruments in Table B.2. Specifically, we instrument for $G^c_{i,t}$ and $G^u_{i,t}$ using the interaction of the change in national spending and state indicators. Relative to our baseline, we do not split national spending by whether it is more- or less-constrained. While the multiplier on constrained spending increases to zero, we continue to find that unconstrained spending has a significantly larger multiplier.

Table B.1: Regional multiplier for total government spending

	GDP	
	Extended Sample	Baseline Sample
Total ($G_{s,t}$)	1.00	1.48
	(0.41)	(0.54)
Observations	1786	1129
Sample	1983–2019	1996–2019

NOTE. This table presents IV estimates of equation (B.1). We instrument for the change in state-level spending, relative to GDP, using the interaction of the analogous change national spending and state indicators. Standard errors clustered by state and year are reported in parentheses.

Table B.2: Alternative instrument: national total government spending

	GDP
Constrained ($G^c_{s,t}$)	-0.06
	(0.93)
Unconstrained ($G^u_{s,t}$)	2.68
	(0.80)
Observations	1129
p(uncon<con)	0.03
Sample	1996–2019

NOTE. This table presents IV estimates of equation (2). We instrument for the change in state-level spending, relative to GDP, using the interaction of the analogous change national spending, state indicators. We test the null hypothesis that the coefficient on constrained spending is larger than the coefficient on unconstrained spending and report the associated p value in the row labeled "p(uncon<con)." Standard errors clustered by state and year are reported in parentheses.

B.2 Industry Composition

Our baseline breakdown of procurement contracts is based on the GPA. There, we say that unconstrained contracts are those covered by the GPA: These are contracts for a certain set of products, Ω, that are above a certain value threshold, Θ. Constrained contracts is the remainder of defense contracts. There is therefore a set of products—those not in Ω—that are only in the constrained contracts. This raises the concern that we may only be identifying the effects of spending on different industries.

We have created another breakdown of contracts that does not depend on which product is being purchased. We instead rely only on the threshold, Θ, to determine whether a contract is constrained or unconstrained. Constrained contracts are those with dollar amounts between zero and Θ. Unconstrained contracts are those with dollar amounts between Θ and 2Θ. The second condition ensures that unconstrained contracts are roughly similar in size to constrained contracts, i.e., it does not include the relatively fat right tail of procurement contracts highlighted by Cox et al. (2024). We use this breakdown to construct the national instruments for equation (2), which ensures that the only variation we use for identification comes from the contract size, and not the product composition.

We present the results in Table B.3. We continue to find statistically and economically larger multipliers for unconstrained spending. Thus, it does not appear that the potential difference industry composition in our baseline breakdown of contracts is driving our results.

Table B.3: Threshold-only breakdown

	GDP
Constrained ($G^c_{s,t}$)	-1.74
	(1.16)
Unconstrained ($G^u_{s,t}$)	2.80
	(1.51)
Observations	1129
p(uncon<con)	0.04
Sample	1996–2019

Note. This table presents IV estimates of equation (2) using our alternative definition of constrained and unconstrained contracts based solely on the GPA threshold, and not on the GPA product list to construct our instruments. Standard errors clustered by state and year are reported in parentheses.

Table B.4: Pre- and post-1996 estimates

	GDP
$G^c_{s,t} \times \mathbf{1}\{y < 1996\}$	0.70
	(0.99)
$G^u_{s,t} \times \mathbf{1}\{y < 1996\}$	1.46
	(1.58)
$G^c_{s,t} \times \mathbf{1}\{y \geq 1996\}$	-0.83
	(0.82)
$G^u_{s,t} \times \mathbf{1}\{y \geq 1996\}$	2.06
	(1.02)
Observations	1786
p(uncon<con), pre-1996	0.31
p(uncon>con), post-1996	0.02
Sample	1983–2019

Note. This table presents IV estimates of equation (2) for different time periods. We interact the constrained and unconstrained state-level spending, and the corresponding national aggregates, by a post-1996 indicator. We test the null hypothesis that the coefficient on constrained spending is larger than the coefficient on unconstrained spending in the post-1996 period, and report the associated p value in the row labeled "p(uncon<con)." We test the null hypothesis that the coefficient on constrained spending is smaller than the coefficient on unconstrained spending in the pre-1996 period, and report the associated p value in the row labeled "p(uncon<con)." Standard errors clustered by state and year are reported in parentheses.

B.3 WTO GPA Placebo

The WTO Government Procurement Agreement definitions that we use to demarcate constrained and unconstrained contracts came into effect in 1996. In Table B.4, we present our baseline specification split by time period. We interact the constrained and unconstrained state-level spending, and the corresponding national aggregates, by a post-1996 indicator. Because the GPA breakdown was not a legislated determinant of the extent of Buy American restrictions before 1996, we can use that period as a placebo test for our treatment variables.

The top two rows of Table B.4 show that the multipliers are similar between constrained and unconstrained contracts before the GPA took effect. Indeed, we cannot reject the null hypothesis that two effects are equal before 1996. After 1996, the effects are similar to our baseline. Therefore, it appears that it is the constraints imposed by the policy, and not some other inherent feature of the contract breakdown, that is driving the result that unconstrained contracts have larger economic effects than constrained contracts.

Table B.5: Regional estimates

	GDP
Constrained ($G^c_{s,t}$)	-1.04
	(5.21)
Unconstrained ($G^u_{s,t}$)	3.89
	(1.45)
Observations	240
p(uncon<con)	0.21
Sample	1996–2019

NOTE. This table presents IV estimates of equation (2) using data aggregated to the regional level, rather than at the state-level. Our definition of regions follows that of Nakamura and Steinsson (2014): It is nearly identical to the regions defined by the Census, except that we split the South Atlantic region into two. One includes Delaware, Maryland, Washington DC, Virginia, and West Virginia. The other includes North Carolina, South Carolina, Georgia, and Florida. Standard errors clustered by region and year are reported in parentheses.

B.4 Regional Estimates

In Table B.5, we present estimates of a version of our baseline specification in which we replace state-level data with region-level data. Under the assumption that region-to-region spillovers are less likely than state-to-state spillovers, this serves as a test of whether the difference in constrained vs. unconstrained multipliers that we identify arises simply from the fact that constrained spending might simply require more inter-U.S. transfers of goods. We find, instead, that the difference in regional multipliers is larger than the difference in state-level multipliers, albeit not statistically significant given the imprecision of the estimate of the constrained spending multiplier. This suggests that it is within-state differences in how procurement contracts are fulfilled that drives the difference in their differential economic effects.

Table B.6: Bartik Instrument

	GDP
Constrained ($G^c_{s,t}$)	-3.59
	(1.67)
Unconstrained ($G^u_{s,t}$)	4.33
	(2.31)
Observations	1129
p(uncon<con)	0.02
Sample	1996–2019

NOTE. This table presents IV estimates of equation (2) using an alternative set of instruments than in our baseline. We instrument for state-level constrained spending growth using the product of national-level constrained spending growth and the average fraction of constrained state-level spending to GDP from 1979–1995. The unconstrained instrument is analogous. Standard errors clustered by state and year are reported in parentheses.

B.5 Bartik Instrument

In Table B.6, we present the results of estimating our baseline specification using Bartik instruments. We construct the instrument by multiplying national constrained spending with the average ratio of state-level constrained spending to GDP between 1979 and 1995. We construct the instrument for unconstrained spending analogously. We find results that are consistent with our baseline: larger, and positive, effects of unconstrained spending.

B.6 Dynamic Responses

Figure B.1 reports the dynamic cross-sectional multiplier regression discussed in Section 5.

Figure B.1: Cross-Sectional Multiplier Local Projection

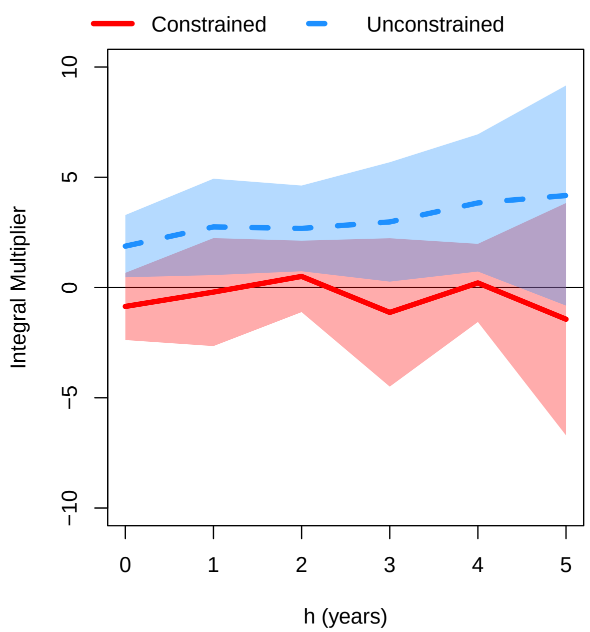

Note. This table presents a dynamic estimates of equation (2), in which we replace the left-hand side with $(Y_{s,t+h} - Y_{s,t-2})/Y_{s,t-2}$.

C Model

C.1 Households

The economy has two regions: a and b. Each region is composed of a continuum of household types indexed by x, where x indicates the type of labor supplied by the household. The measure of agents in regions a and b are n and $1-n$, respectively. Household's in region a maximize

$$\mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t u(C_t, \mathcal{L}_t(x)).$$

where

$$u(C_t, \mathcal{L}_t(x)) = \frac{C_t^{1-\frac{1}{\sigma}}}{1-\frac{1}{\sigma}} - \chi \frac{\mathcal{L}_t(x)^{1+\frac{1}{\nu}}}{1+\frac{1}{\nu}}$$

subject to the flow budget constraint

$$P_t C_t + \mathbb{E}_t[\mathcal{M}_{t,t+1} B_{t+1}(x)] \leq B_t(x) + W_t(x)\mathcal{L}_t(x) + \int_0^1 \Xi_{at}(z) dz - T_t$$

and a transversality condition. The good consumed is a CES aggregate of the goods produced in regions a and b, given by:

$$C_t = \left[\phi_b^{\frac{1}{\eta}} C_{at}^{\frac{\eta-1}{\eta}} + \phi_b^{\frac{1}{\eta}} C_{bt}^{\frac{\eta-1}{\eta}} \right]^{\frac{\eta}{\eta-1}}$$

The good in each region is also a CES composite

$$C_{at} = \left[\int_0^1 c_{at}(z)^{\frac{\theta-1}{\theta}} dz \right]^{\frac{\theta}{\theta-1}} \qquad C_{bt} = \left[\int_0^1 c_{bt}(z)^{\frac{\theta-1}{\theta}} dz \right]^{\frac{\theta}{\theta-1}}$$

Each household allocates labor between production that satisfies the Buy-American Act and production that does not:

$$\mathcal{L}_t(x) = \widetilde{L}_t(x) + L_t(x)$$

It is costless to transition between the two types of production.

Combining the household's optimality conditions for consumption and savings yields the household's Euler equation

$$\mathcal{M}_{t,t+1} = \beta \frac{C_t^{\frac{1}{\sigma}} P_t}{C_{t+1}^{\frac{1}{\sigma}} P_{t+1}}$$

where the stochastic discount factor is defined as

$$\mathcal{M}_{t,t+1} = \frac{\beta \lambda_{t+1}}{\lambda_t}.$$

Combing labor and consumption optimality yields

$$\frac{\chi \mathcal{L}_t(x)^{\frac{1}{\nu}}}{W_t(x)} = \frac{1}{P_t C_t^{\frac{1}{\sigma}}}.$$

Households optimally cost-minimize over varieties in order to attain C_t. This yields the following demand curves for goods from a and b

$$C_{at} = \phi_a C_t \left(\frac{P_{at}}{P_t} \right)^{-\eta} \qquad C_{bt} = \phi_b C_t \left(\frac{P_{bt}}{P_t} \right)^{-\eta}$$

And between varieties produced in a and b

$$c_{at}(z) = C_{at} \left(\frac{p_{at}}{P_{at}} \right)^{-\theta} \qquad c_{bt}(z) = C_{bt} \left(\frac{p_{bt}}{P_{bt}} \right)^{-\theta}$$

where the price indexes and a and b are

$$P_{at} = \left[\int_0^1 p_{at}^{1-\theta} dz \right]^{\frac{1}{1-\theta}} \qquad P_{bt} = \left[\int_0^1 p_{bt}^{1-\theta} dz \right]^{\frac{1}{1-\theta}}$$

and the aggregate price level in a:

$$P_t = \left[\phi_a P_{at}^{1-\eta} + \phi_b P_{bt}^{1-\eta} \right]^{\frac{1}{1-\eta}}.$$

Region b is analogous to region a. We denote variables in region b with asterisks. The corresponding equilibrium conditions are

$$P_t^* C_t^* + \mathbb{E}_t[\Lambda_{t,t+1} B_{t+1}^*(x)] \leq B_t^*(x) + W_t^*(x)\mathcal{L}_t^*(x) + \int_0^1 \Xi_{at}^*(z)dz - T_t^* \tag{C.1a}$$

$$\Lambda_{t,t+1} = \beta \frac{(C_t^*)^{\frac{1}{\sigma}} P_t^*}{(C_{t+1}^*)^{\frac{1}{\sigma}} P_{t+1}^*} = \beta \frac{(C_t^*)^{\frac{1}{\sigma}} P_t Q_t}{(C_{t+1}^*)^{\frac{1}{\sigma}} P_{t+1} Q_{t+1}} \tag{C.1b}$$

$$\frac{\chi(\mathcal{L}_t^*(x))^{\frac{1}{\nu}}}{W_t^*(x)} = \frac{1}{P_t^*(C_t^*)^{\frac{1}{\sigma}}} \tag{C.1c}$$

$$C_{at}^* = \phi_a^* C_t^* \left(\frac{P_{at}}{P_t^*}\right)^{-\eta} = \phi_a^* C_t^* \left(\frac{P_{at}}{P_t Q_t}\right)^{-\eta} \tag{C.1d}$$

$$C_{bt}^* = \phi_b^* C_t^* \left(\frac{P_{bt}}{P_t^*}\right)^{-\eta} = \phi_b^* C_t^* \left(\frac{P_{bt}}{P_t Q_t}\right)^{-\eta} \tag{C.1e}$$

$$c_{at}^*(z) = C_{at}^* \left(\frac{p_{at}(z)}{P_{at}}\right)^{-\theta} \tag{C.1f}$$

$$c_{bt}^*(z) = C_{bt}^* \left(\frac{p_{bt}(z)}{P_{bt}}\right)^{-\theta} \tag{C.1g}$$

$$P_t^* = \left[\phi_a^* P_{at}^{1-\eta} + \phi_b^* P_{bt}^{1-\eta}\right]^{\frac{1}{1-\eta}} \tag{C.1h}$$

where $Q_t \equiv \frac{P_t}{P_t^*}$ is the real exchange rate.

C.2 Government

The government can purchase goods according to the private sector technology or the Buy-American-specific technology, indicated with tildes, from region a or region b. We denote these four types of per-capita government spending by:

$$G_{at}, \widetilde{G}_{at}, G_{bt}, \widetilde{G}_{bt}$$

Each follows an exogenous AR(1) process.

The government has the same preferences among varieties of each type of good as households. The demand curves are therefore:

$$g_{at}(z) = G_{at}\left(\frac{p_{at}(z)}{P_{at}}\right)^{-\theta} \qquad \widetilde{g}_{at}(z) = \widetilde{G}_{at}\left(\frac{\widetilde{p}_{at}(z)}{\widetilde{P}_{at}}\right)^{-\theta} \tag{C.2a}$$

$$g_{bt}(z) = G_{bt}\left(\frac{p_{bt}(z)}{P_{bt}}\right)^{-\theta} \qquad \widetilde{g}_{bt}(z) = \widetilde{G}_{bt}\left(\frac{\widetilde{p}_{bt}(z)}{\widetilde{P}_{bt}}\right)^{-\theta} \tag{C.2b}$$

The government runs a balanced budget financed by lump-sum taxes

$$n[P_{at}G_{at} + \widetilde{P}_{at}\widetilde{G}_{at}] + (1-n)[P_{bt}G_{bt} + \widetilde{P}_{bt}\widetilde{G}_{bt}] = T_t. \tag{C.2c}$$

The central bank sets the nominal interest in response to CPI inflation and GDP

$$(1+i_t) = (1+i_{t-1})^{\rho_r}\left(\mathbf{\Pi}_t^{\phi_\pi}\mathbf{Y}_t^{\phi_\pi}\right)^{1-\rho_r} \tag{C.3a}$$

where

$$\mathbf{\Pi}_t \equiv \Pi_t^n(\Pi_t^*)^{1-n} \qquad\qquad \mathbf{Y}_t \equiv \mathcal{Y}_t^n(\mathcal{Y}_t^*)^{1-n}$$

The interest rate is the return to a one-period, riskless nominal bond, i.e.

$$\mathbb{E}_t[M_{t,t+1}(1+i_t)] = 1. \tag{C.3b}$$

C.3 Firms

There is a continuum of firms indexed by z in region a. Firm z specializes in the production of differentiated good z to produce output $y_{at}(z)$. The production function of each firm z is:

$$y_{at}(z) = \phi\left[\alpha L_t(z)^\rho + (1-\alpha)M_t(z)^\rho\right]^{\frac{1}{\rho}}$$

where $M_t(z)$ is an imported intermediate input. It is supplied inelastically at price μ by the rest of the world.

The firm also produces for the government using the technology

$$\widetilde{y}_{at}(z) = \phi\left[\widetilde{\alpha}\widetilde{L}_t(z)^\rho + (1-\widetilde{\alpha})\widetilde{M}_{,t}(z)^\rho\right]^{\frac{1}{\rho}}.$$

We assume labor is mobile across sectors, but immobile across regions. The associated unit cost of production of the private-sector good is given by:

$$S_t(z) = \left[\alpha^{\frac{1}{1-\rho}}w_t(x)^{\frac{\rho}{\rho-1}} + (1-\alpha)^{\frac{1}{1-\rho}}\mu^{\frac{\rho}{\rho-1}}\right]^{\frac{\rho-1}{\rho}}$$

and the unit cost for the BAA good is

$$\widetilde{S}_t(z) = \left[\widetilde{\alpha}^\sigma w_t(x)^{1-\sigma} + (1-\widetilde{\alpha})^\sigma(\tau\mu)^{1-\sigma}\right]^{\frac{1}{1-\sigma}}.$$

Firm z maximizes lifetime expected profits:

$$E_t \sum_{j=0}^{\infty} \Lambda_{t,t+j} \left[p_{at+j}(z) y_{at+j}(z) + \widetilde{p}_{at+j}(z) g_{at}(z) - w_{t+j}(x)(L_{t+j}(z) + \widetilde{L}_{t+j}(z)) - p_{M,t+j}(M_t + \widetilde{M}_t)_{t+j}(z) \right] \tag{C.4}$$

where $\Lambda_{t,t+j}$ is the stochastic discount factor.

Firms take the industry wage, $w_t(x)$ and the price of imported intermediates, μ_t as given. Optimal choices of labor and imported input demand for unconstrained and constrained production are given by:

$$L_t(z) = \frac{y_{at}(z)}{\phi} \left[\frac{\alpha \phi S_t(z)}{w_t(x)} \right]^{\frac{1}{1-\rho}} \qquad \text{and} \qquad \widetilde{L}_t(z) = \frac{\widetilde{y}_{at}(z)}{\phi} \left[\frac{\alpha \phi \widetilde{S}_t(z)}{w_t(x)} \right]^{\frac{1}{1-\rho}}$$

and

$$M_t(z) = \frac{y_{at}(z)}{\phi} \left[\frac{(1-\alpha) \phi S_t(z)}{\mu} \right]^{\frac{1}{1-\rho}} \qquad \text{and} \qquad \widetilde{M}_t(z) = \frac{\widetilde{y}_{at}(z)}{\phi} \left[\frac{(1-\alpha) \phi \widetilde{S}_t(z)}{\tau \mu} \right]^{\frac{1}{1-\rho}}$$

Firm z can reoptimize its price with probability $1 - \kappa$. With probability κ it must keep its price unchanged. Optimal price setting by firm z in periods where it can change its price implies:

$$p_{at}(z) = \frac{\theta}{\theta - 1} \mathbb{E}_t \sum_{j=0}^{\infty} \frac{\kappa^j \Lambda_{t,t+j} y_{at+j}(z)}{\mathbb{E}_t \sum_{k=0}^{\infty} \kappa^k \Lambda_{t,t+k} y_{at+j}(z)} S_{t+j}(z)$$

for unconstrained spending, and

$$\widetilde{p}_{at}(z) = \frac{\theta}{\theta - 1} \mathbb{E}_t \sum_{j=0}^{\infty} \frac{\kappa^j \Lambda_{t,t+j} \widetilde{y}_{at+j}(z)}{\mathbb{E}_t \sum_{k=0}^{\infty} \kappa^k \Lambda_{t,t+k} \widetilde{y}_{at+j}(z)} \widetilde{S}_{t+j}(z)$$

for constrained spending.

C.4 Market Clearing

Firm z must produce enough to satisfy demand for its product, which comes from three sources: home consumers, foreign consumers, and the government. Government demand takes two forms: unconstrained demand, which is analogous to private demand, and constrained demand. Regular (unconstrained) production must meet demand, and the same goes for constrained production, yielding the following conditions:

$$y_{at}(z) = (nC_{at} + (1-n)C^*_{at} + nG_{at}) \left(\frac{p_{at}(z)}{P_{at}}\right)^{-\theta} \tag{C.5a}$$

$$\widetilde{y}_{at}(z) = n\widetilde{G}_{at} \left(\frac{\widetilde{p}_{at}(z)}{\widetilde{P}_{at}}\right)^{-\widetilde{\theta}} \tag{C.5b}$$

$$y_{bt}(z) = (nC_{bt} + (1-n)C^*_{bt} + (1-n)G_{bt}) \left(\frac{p_{bt}(z)}{P_{bt}}\right)^{-\theta} \tag{C.5c}$$

$$\widetilde{y}_{bt}(z) = (1-n)\widetilde{G}_{bt} \left(\frac{\widetilde{p}_{bt}(z)}{\widetilde{P}_{bt}}\right)^{-\widetilde{\theta}} \tag{C.5d}$$

Equating supply and demand the labor market yields

$$\mathcal{L}_t(z) = L_t(z) + \widetilde{L}_t(z) \tag{C.5e}$$

$$\mathcal{L}^*_t(z) = L^*_t(z) + \widetilde{L}^*_t(z) \tag{C.5f}$$

Finally, we impose the risk-sharing condition that the stochastic discount factor is equal in the two regions, which implies that

$$\left(\frac{C^*_t}{C_t}\right)^{-\frac{1}{\sigma}} = \frac{P^*_t}{P_t} \equiv Q_t \tag{C.5g}$$

The model's equilibrium paths for the endogenous variables are described by the optimality conditions and constraints of households in regions a and b; the government's demand and budget constraint (equations (C.2)); the monetary policy rule (equation (C.3a)); and the optimality conditions and constraints of the firms in regions a and b. We also define real GDP per capita as

$$\mathcal{Y}_t = \frac{P_{at}\left(Y_{at} + \widetilde{Y}_{at}\right) - \mu\left(M_t + \tau\widetilde{M}_t\right)}{P_{at}}.$$

C.5 Solution

We log-linearize the model's equilibrium conditions around a zero-inflation steady-state. The process is standard and tedious, so we omit the details. We solve the log-linearized model using the methods of Sims (2002).