NBER WORKING PAPER SERIES

TRADE PROTECTION, STOCK-MARKET RETURNS, AND WELFARE

Mary Amiti
Sang Hoon Kong
David Weinstein

Working Paper 28758
http://www.nber.org/papers/w28758

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
May 2021, Revised August 2022

This paper supersedes an earlier unpublished paper, “The Effect of the U.S.-China Trade War on U.S. Investment.” We thank George Alessandria, Kirill Borusyak, Dave Donaldson, Patrick Farrell, Ethan Ilzetzki, Peter Neary, Michael Peters, Ashwin Rao, Steven Redding, Peter Schott, Yeji Sung, James Tybout, Jon Vogel, and Michael Waugh for many useful comments. We thank Richard Crump for sharing inflation expectations data and Daniel Lewis for sharing the economic surprise data. We also thank Yingjie (Angela) Wu and Dongcheng Yang for providing excellent research assistance. The views expressed in this paper are those of the authors and do not necessarily reflect the position of the Federal Reserve Bank of New York, the Federal Reserve System, or the National Bureau of Economic Research.

Trade Protection, Stock-Market Returns, and Welfare
Mary Amiti, Sang Hoon Kong, and David Weinstein
NBER Working Paper No. 28758
May 2021, Revised August 2022
JEL No. F13,F14,F16

ABSTRACT

This paper develops a method of using financial data to understand the expected impact of trade policy on welfare and other real variables when the policy has heterogeneous and uncertain impacts on firms. We embed a firm-level specific factors model in a consumption capital asset pricing model to map expected cash-flow movements into expected movements in productivity, wages, and welfare. Using our framework, we find that the U.S.-China trade war caused increases in uncertainty and large declines in U.S. stock prices, expected cash flows, and expected productivity. We estimate that the expected decline in U.S. welfare is 4.9 percentage points.

Mary Amiti
Microeconomic Studies Function
Federal Reserve Bank of New York
33 Liberty St
New York, NY 10045
and CEPR
mary.amiti@gmail.com

Sang Hoon Kong
Columbia University
sanghoonkong@gmail.com

David Weinstein
Columbia University
Department of Economics
420 W. 118th Street
MC 3308
New York, NY 10027
and NBER
dew35@columbia.edu

A data appendix is available at http://www.nber.org/data-appendix/w28758

1 Introduction

This paper develops a method of using financial data to understand the expected impact of trade policy on welfare and other real variables when the policy has heterogeneous and uncertain impacts on firms. Our approach improves on existing methodologies in several ways. First, relative to the calibration of canonical static trade models, our approach does not require researchers to assume that the policy has no effects on firm-level productivity. Second, although dynamic models do allow for changes in firm-level productivity, calibrations of these models typically require researchers to assume uniform tariff reductions and a functional form for productivity spillovers.[1] By contrast our approach allows for arbitrary patterns of tariff reductions and productivity effects. Third, we improve on standard stock-market event studies by developing a methodology for using financial data to estimate general equilibrium effects of a policy announcement. Financial data is particularly well suited for understanding the impact of protection in general equilibrium because markets are forward looking and agents have strong incentives to incorporate even complex dynamics and productivity effects into asset prices. We apply this methodology to the U.S.-China trade war and find impacts of the trade war that are an order of magnitude larger than those typically found in static studies that allow for no productivity effects.

Our method requires us to solve two theoretical challenges. The first is to specify a general equilibrium production structure that can be integrated into an asset-pricing model. We show how to adapt Jones (1975)'s *industry* specific factors model into a *firm-level* specific factors model to describe the production side of the economy. In this setup, payments to firm-specific factors equal firm cash flows (i.e., revenues less variable costs). A key difference is that the Jones model maps price changes into output, employment, and factor prices, whereas we invert this logic to show how the knowledge of returns to the specific factor can be used to identify the other variables Jones considers. This difference is important because it avoids having to specify how tariffs affect all prices in the economy and instead only requires knowledge of how the tariff affects the returns to the specific factor. We show that the returns to the specific factor are the same as the change in firm cash flow in the firm-level specific factors model. This insight enables us to derive analytic solutions for how movements in expected cash flow map into movements in expected firm-level sales, wages, employment, markups, and both quantity and revenue total factor productivity (TFPR). Because the mapping between each of these variables and cash flow is linear, a researcher with an estimate of how a policy affected expected firm cash flow can compute how it must affect the expected value of these variables.

The close relationship between expected cash-flow movements and TFPR highlights an important theoretical result that makes the specific factors model a particularly good framework for understanding tariffs' effects on productivity. The central explanatory variable in Jones (1975)'s specific factors model is the change in the effective rate of protection (ERP), which he defined as the log change in output prices less a weighted average of input price changes, all divided by the share of value added in sales. We show that log changes in ERP equal the dual measure of the change in TFPR. Thus, if tariffs lower

[1] Examples of prominent dynamic models include Sampson (2016), Buera and Oberfield (2020), and Perla et al. (2021).

ERP, they must drive down TFPR. The intuition for this result is that primal TFPR is the residual between revenue growth less a weighted average of input growth while ERP is the residual between output price growth and a weighted average of input price changes. Thus, they are two ways of measuring the same object.

The second theoretical challenge is to integrate the production structure of the specific factors model into one that will allow us to model the welfare impacts of the policy. Doing so requires us to have a welfare function and a rigorous method of mapping movements in financial variables—in our case stock and bond prices—into expected cash flows and prices. We do this by building off the standard consumption capital asset pricing model (CCAPM). We modify the CCAPM to be a two-period version in which the production structure of the economy in each period is given by the specific factors model. This simplification enables us to specify the impact of policies in terms of the present discounted value of their effects on cash flow, which enables us to avoid having to specify the time path of the policy impacts.

The CCAPM framework also enables us to measure how a policy affects welfare through two channels: direct impacts on consumption growth and changes in agents' perceptions of consumption uncertainty. Our model assumes investors hold shares in firms in order to obtain access to future firm cash flows, but their willingness to pay for shares depends on the stochastic discount factor (SDF), which is a function of expected future consumption levels and uncertainty about future consumption. This structure has two useful properties. First, it provides a link from observed stock price movements to expected cash flow and consumption changes. Second, if shocks to prices are log-normally distributed, we show how the policy's welfare impact can be calculated based on its impact on expected firm cash flows, tariff revenues, inflation, and consumption volatility.

Our empirical specification builds off the CCAPM to show how to use data on stock returns and policy announcement dates to identify the policy impact on expected firm cash flows. There are several steps in this procedure. First, we use a factor model to identify a policy's impact on stock returns due to its effect on latent macro variables. Second, we use the approach of Campbell and Vuolteenaho (2004) to separate out the movements in the latent variables that are due to shifts in discount rates and those that are due to cash flow. This procedure lets us estimate what we term the "macro effect" of the policy: the movement in the expected present discounted value of firm cash flow arising from the policy's expected impact on macro variables. For example, if a tariff announcement causes an exchange rate appreciation that affects firm cash flows, we would consider this part of the macro effect. Third, we recognize that the residual from this estimation equation is the abnormal return, which captures differential changes in stock returns due changes in expected cash flow. If we use this residual as the dependent variable in a standard stock-market event study in which we use firm exposure variables (e.g., whether a firm imports from China) as explanatory variables, we can estimate the treatment effects of the policy. These treatment effects correspond to the policy's expected *differential* impact on the present discounted value of firm cash flow. Summing together the macro and treatment effects gives us our estimate of the expected impact of the policy on the present discounted value of cash flow.

To calculate the welfare impact of the policy, we also need to know its impact on expected tariff revenue, inflation, and consumption volatility. We obtain our estimates of the policy's impact on expected tariff revenues from Fajgelbaum et al. (2020). We identify the effect of the tariff announcements on expected inflation by movements in expected inflation estimates derived from bond prices during event windows. Our theoretical framework allows us to calculate the policy's impact on expected consumption from the estimated effects on expected firm cash flows, tariff revenue, and inflation. Lastly, we identify the policy's impact on consumption volatility by comparing the variance in daily changes to expected consumption implied by stock-market data during event windows to the variance during the days preceding the event windows.

We apply this framework to understanding the implications of the U.S.-China trade war. In order to identify a trade-war announcement, we search for the first mention in the media of the tariffs implemented by the U.S. or China during the 2018-2019 trade war, resulting in eleven trade war events.[2] We find that the U.S.-China trade-war announcements are associated with large stock-price declines regardless of whether we look at impacts over one-day, three-day, or five-day event windows. We find that during a three-day window around each of the trade-war announcements, stock prices fell 12.9 percent in total: a $3.7 trillion loss in market value. These movements imply that markets viewed the trade war announcements as having significant impacts on consumption, uncertainty, and/or cash flow.

Before measuring the welfare effects, we provide external validation of our estimates of expected cash flow by showing that they predict movements in real outcome variables as predicted by theory. We first test whether expected cash-flow changes derived from financial data during tariff event windows predict future *accounting* cash-flow movements and does not cause movements in firm-specific factors, which we measure by intangible assets. We find that these predictions are borne out in the data. Second, our theory allows us to derive relationships between expected cash flow and expected employment, output, tangible asset, and productivity movements. All of these variables move in the ways that our theory predicts.

Next, we turn to measuring their implied impact on welfare. Consistent with the results of Campbell and Vuolteenaho (2004), we find that a large share of the drop in stock prices on tariff-announcement days was due to shifts in discount rates. However, we still find that movements in cash flow account for 7.0 percentage points of the 12.9 percentage points fall in stock prices. In addition, we find that our estimates of expected consumption volatility doubled over the course of the eleven events. Jointly, we find that these lowered U.S. welfare by 4.9 percent. Of this decline, virtually all of it comes from lower expected consumption, with increased consumption uncertainty explaining only a trivial amount of the fall. The small consumption uncertainty effect stems from the fact that the variance of consumption in the model and in aggregate US data is quite small, so doubling its magnitude does not harm consumers much.

Our model and empirics suggest that there are two important channels explaining the large welfare impact of tariffs in our study relative to those obtained from static models: productivity and macro effects. First, our baseline results show that the drop in stock prices implies that market participants expect future U.S. TFP to be 4.5 percent lower because of the trade war than would otherwise be expected. This large TFP effect helps

[2] We used Factiva and Google search to identify the event dates.

explain why we estimate expected real wages to fall by 4.3 percent as a result of the tariffs. Second, static models only allow tariffs to affect the relative prices, and therefore do not allow tariffs to affect macro variables. In our study, these treatment effects account for less than one quarter of the total effect. Most of the impact comes through the macro term, which captures forces like exchange-rate changes and more general productivity spillovers.

Related Literature Our work is closely related to the voluminous literature on stock-market event studies that use trade data (Grossman and Levinsohn (1989), Hartigan et al. (1986), Breinlich (2014), Fisman et al. (2014), Moser and Rose (2014), Breinlich et al. (2018), Crowley et al. (2019), Huang et al. (2019), and Greenland et al. (2020)). We differ in the use of a general equilibrium model to interpret the data. Greenland et al. (2020) is particularly relevant in that they show that positive firm abnormal returns in response to lower trade uncertainty, through the granting of permanent normal trade relations in 2000, led to future increases in firm employment, sales, productivity and profits. Our approach yields a theoretical foundation for their regressions, and their results validate our assumption that movements in expected cash flows are tightly linked to movements in future accounting profits and other non-financial variables. We also document a significant link between firm stock returns and future movements in non-financial variables using a structural approach to measuring the impact of policy announcements.

The specific factors model, which forms the basis of our approach, has also been used extensively in empirical estimation in recent years (c.f., Topalova (2010), Kovak (2013), and Dix-Carneiro and Kovak (2017)). These papers have shown that many of the large effects of trade policy changes on wages often take a decade to be fully apparent in the data. Our paper provides a complementary way of thinking about the long-term effects of a policy change in terms of expected wages. In particular, most papers in this literature only look at the impact of output tariffs, so tariffs are assumed to always raise the effective rate of protection. However, in our setup, we allow tariffs to affect input prices as well, so the imposition of tariffs can either raise the ERP by increasing firm output prices or lower it by raising the cost of the firm's imported intermediate inputs.

Our paper is related to the vast empirical trade literature over the last two decades showing that trade liberalizations have big effects on per capita income and productivity. These studies have shown that *firm-level* TFP is very sensitive to ERP and import competition more generally.[3] We also identify large impacts of trade policy on revenue TFP, but our identification is based on using stock-price data filtered through a general equilibrium model. Our paper is also related to the macro literature evaluating the impact of trade on income that has found evidence of large impacts of trade on productivity and income (c.f., Frankel and Romer (1999); Alcalá and Ciccone (2004); Feyrer (2019)). These studies find that the elasticity of per capita income with respect to trade ranges from 0.5 to 3 and that most of the effect arises through trade's impact on productivity. Although our

[3]For example, Amiti and Konings (2007) estimate the elasticity of firm-level TFP with respect to input tariffs to be -1.2 in Indonesia for firms that import their inputs. There were also gains to non-importers but these were smaller, so the average elasticity across all firms was -0.44. Topalova and Khandelwal (2011) estimate the elasticity to be -0.5 in Indian data, and Brandt et al. (2017) and Brandt et al. (2019) estimate the elasticity to be -2.3 in Chinese data. Bloom et al. (2016) find that Chinese import competition accounts for 14% of European technology upgrading.

work also finds large impacts of trade on productivity and welfare, an important difference between our work and the macro literature is that we build these estimates up from firm-level data on stock prices and use a structural general equilibrium setup to obtain our estimates.

We also contribute to the burgeoning literature on understanding the importance of protection for the economy through macro or policy uncertainty channels (Baker et al. (2016); Pierce and Schott (2016); Handley and Limão (2017); Caldara et al. (2019); Greenland et al. (2020)). Like these papers, our paper also suggests that trade policy announcements can have impacts that arise through uncertainty or changing the macro environment, but we differ in our use of financial data to identify the shocks and the use of a general equilibrium model. Our paper is also related to work on the China shock. For example, Autor et al. (2013) and Caliendo et al. (2019) show how trade with China affected U.S. employment, wages, and welfare, but our work focuses on trade policy announcements.

Finally, our paper is related to the literature documenting the impact of the trade war on prices (c.f., Amiti et al. (2020); Fajgelbaum et al. (2020); Flaaen et al. (2020); Amiti et al. (2019); Cavallo et al. (2021)). These papers have found that during the U.S.-China trade war, tariff passthrough into import prices was close to complete, consistent with our finding that higher U.S. tariffs negatively affected importers. Cavallo et al. (2021) found that Chinese tariffs depressed U.S. exporter prices, also consistent with our findings of negative abnormal returns for firms exporting to China following Chinese retaliation events.

2 Theory

This section develops a general-equilibrium model that provides a mapping from expected movements in firm cash flows into expected movements in returns to specific factors, wages, employment, output, and welfare. We build the theory in a number of steps in order to show how one can embed the production structure of a static specific factors model into the CCAPM. Our ultimate model is composed of three elements—a specific factors model of goods and factor prices, a method of integrating stochastic price shocks into the setup, and a model of consumption and asset pricing—so the sections of our theory develop and integrate each of these elements into our setup.

First, in Section 2.1, we develop a two-period, firm-level, specific factors model of production and show that movements in firm cash flow (which are identical to the returns to the specific factor in this setup) are sufficient statistics to infer the movements in wages, firm sales, employment, prices, markups, ERP and TFP. This structure enables us to do comparative static exercises in the absence of asset markets or uncertainty about future prices. We then introduce uncertainty into this setup letting prices be stochastic in period two. We show that our baseline results about how equilibrium variables are related to one another holds in expectation, i.e., the mapping between expected cash-flow movements and each expected outcome variable is identical to those in the certainty case.

Second, in Section 2.2, we embed the specific factors model into a two-period CCAPM in which the value of holding a firm's share today is determined by the investor's expectations about the SDF, inflation, and firm's cash flow tomorrow. The CCAPM structure enables us to write expected consumption as a function of expected cash flow, tariff revenues, and consumer prices. It also allows us to obtain an expression for welfare as a function of how a policy affects expected consumption and the variance in consumption.

2.1 Production

Initially, we abstract from any uncertainty about prices and show that if we invert the Jones (1975) model we can write wages, firm-level labor and output changes as a function of movements in returns to specific factors. We then introduce price uncertainty in Section 2.1.2.

2.1.1 Production with Price Certainty

We assume that there is a fixed number of firms in the economy indexed by $f = 1, ..., F$. Each firm owns a fixed, non-tradable specific factor, V_f. The firm uses a constant returns to scale technology to produce y_f units of output by combining its firm-specific factor V_f, with L_f homogeneous workers hired in competitive labor markets, and differentiated intermediate inputs indexed by i. We denote the unit cost of production as $c_f(w, r_f, q_1, ..., q_n)$, where the arguments correspond to the wage (w), the shadow price of the intangible asset (r_f), and prices of a set of intermediate inputs $(q_1, ..., q_n)$. Firms maximize their cash flow (profits) defined as their sales less variable costs, which are then paid out to the firm's owners. That is, $r_f V_f = (p_f - c_f(w, 0, q_1, ..., q_n)) y_f$, where p_f is the price of firm f's output.

There are two noteworthy features of this setup. First, the payment received by the firm's owners $(r_f V_f)$ is identical to the firm's cash flow, and we will therefore call this term cash flow for the remainder of the paper. Second, because the amount of each firm's specific asset is fixed ($\hat{V}_f = 0$), log changes in the shadow price of the specific factor equal the log change in firm cash flow (i.e., $\hat{r}_f = \widehat{r_f V_f}$). Thus, we will henceforth refer to $\hat{r}_f$ as the log change in the firm's cash flow.

As in Jones (1975), we impose the full-employment conditions on labor and each firm's specific factor:

$$\sum_f a_{Lf} y_f = L, \text{ and} \tag{1}$$

$$a_{Vf} y_f = V_f, \tag{2}$$

where $L \equiv \sum_f L_f$ and the unit input requirements for labor and the specific factor are given by a_{Lf} and a_{Vf}. Since $a_{Lf} y_f = L_f$, the first full-employment condition (1) stipulates that firm-level employment will adjust with firm-level production. In contrast, the second full-employment condition (2) stipulates that the unit-input requirement of the specific factor (a_{Vf}) is inversely proportional to firm output (y_f) because the amount of the specific factor (V_f) is fixed. Similarly, we assume a common, constant elasticity of substitution between the specific factor and labor (σ), so that the factor intensity of production (a_{Vf}/a_{Lf}) is determined by relative factor prices and this elasticity:

$$\hat{a}_{Vf} - \hat{a}_{Lf} = \sigma(\hat{w} - \hat{r}_f), \tag{3}$$

where hats over variables indicate log changes. Cost minimization implies that the unit-input requirements can be written as $a_{Lf} = \frac{\partial c_f}{\partial w}$, $a_{Vf} = \frac{\partial c_f}{\partial r_f}$, and $a_{if} = \frac{\partial c_f}{\partial q_i}$, so we have

$$a_{Lf}w + a_{Vf}r_f + \sum_i a_{if}q_i = p_f. \tag{4}$$

We are now ready to prove our first proposition linking changes in cash flows to wages.

Proposition 1. *If the elasticity of substitution between labor and the specific factor for all firms is constant, the log change in wages equals the employment-share weighted average of the log changes in cash flow, i.e.,*

$$\hat{w} = \sum_f \frac{L_f}{L}\hat{r}_f,$$

and the log change in employment in each firm equals $\hat{L}_f = \sigma\left(\hat{r}_f - \sum_{f'} \frac{L_{f'}}{L}\hat{r}_{f'}\right)$.

Proof. See Appendix A.1 □

The intuition behind the first equation in Proposition 1 is that the full-employment condition implies that changes in factor returns cannot yield an increase in the aggregate demand for labor. However, the aggregate demand for labor will only remain constant if the changes in relative wages ($\hat{w} - \hat{r}_f$) are zero "on average," i.e., changes in wages ($\hat{w}$) equal a firm-size weighted change in the average of log changes in cash flow ($\sum_f \frac{L_f}{L}\hat{r}_f$). The second line follows immediately from this equation and the fact that the amount of the specific factor is fixed, so the left-hand side of equation (3) is just $-\hat{L}_f$.[4]

Proposition 1 is based on the structure of Jones (1975) but differs in a number of respects. First, Jones was concerned about mappings from changes in product prices into factor prices. Here, we invert the logic in Jones to show that knowing the log changes in cash flow pins down changes in wages and employment. Second, by assuming that there is one elasticity of substitution between labor and the specific factor, we simplify the expressions in his canonical model and are able to construct a sufficient statistic for computing wage and employment changes using only information on changes in cash flow.[5] Wages move one for one with the employment-weighted average of log changes in cash flow.[6]

The remaining propositions require one additional assumption: namely that the share of expenditures on total intermediate inputs (i.e., the sum of imported and domestic) are a constant fraction of sales.[7] We define ω_{Lf}, ω_{Vf}, and ω_{if} as the expenditures of firm f on

[4]We relax the assumption of a vertical labor supply curve in Appendix A.1.1. Allowing aggregate employment to move with changes in cash flow does not undermine the basic result that we can express equilibrium wage changes as a linear function of changes in cash flow.

[5]By contrast, implementing the Jones approach would require us to know the full set of firm-level elasticities. While the assumption of a single elasticity of substitution is more restrictive, other studies have often adopted even more restrictive assumptions, e.g., assuming that $\sigma = 1$ (c.f., Kovak (2013)). Knoblach and Stöckl (2020) conduct a meta-analysis of 49 studies and find that the value of σ typically falls between 0.4 and 0.7.

[6]At first, it may seem surprising that wages rise one for one with average log changes in cash flow, however, this result is present in other models in which firms have positive operating profits. For example, in Melitz (2003), both per worker welfare and average firm profits are monotonically rising in average productivity.

[7]This is a standard assumption whenever one wants to analyze a value-added production function and is common in the macro literature whenever TFP is defined as the residual from subtracting capital and labor input growth from output growth (see, for example, Hsieh and Klenow (2009)).

labor, the specific factor, and input i expressed as a share of total revenue. We can then obtain an expression for a mapping between relative returns to the specific factor and output changes.

Proposition 2. *If the expenditures on intermediate inputs are a constant fraction of sales, the log change in output is a linear combination of the log changes in returns to the specific factors:*

$$\hat{y}_f = \frac{\omega_{Lf}\sigma}{\omega_{Lf} + \omega_{Vf}} \left(\hat{r}_f - \sum_{f'} \frac{L_{f'}}{L} \hat{r}_{f'} \right)$$

where ω_{Lf} and ω_{Vf} denote the payments to labor and specific factors as a share of revenue.

Proof. See Appendix A.2 □

The assumption that intermediate input expenditures are a constant fraction of sales will be satisfied if production is multiplicatively separable into a composite intermediate input and other factors of production. A Cobb-Douglas production function would satisfy this, but one could also have richer production technologies in which there is an arbitrary elasticity of substitution between labor and firm-specific factors and arbitrary elasticities among intermediate inputs, so long as the elasticity of substitution between the composite intermediate and the other factors of production is one.[8]

We can also use the structure of our model to obtain mappings from cash-flow movements into many other variables of interest. Our starting point is the firm-level definition of the effective rate of protection (ERP):

$$\hat{p}^e_f \equiv \frac{\hat{p}_f - \sum_i \omega_{if}\hat{q}_i}{1 - \sum_i \omega_{if}}. \tag{5}$$

The numerator in this definition is the change in the firm's output prices, less a weighted average of all of the input prices, while the denominator is the share of value added in sales. While Jones (1975) was principally concerned with how movements in firm *prices* (and hence ERP) affect factor prices, a major limitation of his approach is that it is impossible to rigorously map *tariff changes* into ERP changes without making implausible assumptions.[9]

[8]Importantly, this assumption does not imply that the elasticity of substitution between imports and labor is one. For example, suppose that production is given by a Cobb Douglas function: $Y_f = V_f^{\alpha_1} L_f^{\alpha_2} Q_f^{(1-\alpha_1-\alpha_2)}$, where Y_f is output of firm f; Q_f is a composite intermediate used by the firm; and the α_i are parameters between zero and one that satisfy $\alpha_1 + \alpha_2 < 1$. In this case, the elasticity of substitution between the composite input and labor is one. If the composite intermediate input is a function of domestic and imported intermediates (D_f and I_f), so $Q_f = g(D_f, I_f)$, the elasticity of substitution between labor and imported intermediates, could take on values greater than one. For example, if domestic and imported inputs are highly substitutable, the elasticity of substitution between labor and imported intermediates will also be high because the ratio of imported intermediates to labor will fall rapidly when the price of imports rises.

[9]Examples of common implausible assumptions include: firms use no intermediate inputs, perfect or constant passthrough of tariffs into prices, no heterogeneity in firm-level input-output matrices, no effects of tariffs on exchange rates, no impact of tariffs on productivity, etc. For example, a common approach to measuring ERP is to follow Corden (1966) and define it as the change in the output tariff less an input-share

An advantage of our approach is that we can use movements in cash flow to conduct comparative statics exercises and measure a policy's impact on many variables, including ERP, without making any assumptions about how tariffs affect ERP. We proceed by first recalling a result from Jones (1975), who proved that the movement in the returns to each specific factor (i.e., changes in cash flow) can be written as

$$\hat{r}_f = \left(\varphi_f + \frac{1}{\theta_{Vf}} \sum_{f' \neq f} \varphi_{f'}\right) \hat{p}^e_f - \frac{\theta_{Lf}}{\theta_{Vf}} \sum_{f' \neq f} \varphi_{f'} \hat{p}^e_{f'} \text{ and } \hat{w} = \sum_f \varphi_f \hat{p}^e_f, \tag{6}$$

where

$$\varphi_f \equiv \frac{L_f}{\theta_{Vf}} / \sum_{f'} \frac{L_{f'}}{\theta^V_{f'}},$$

θ_{Lf} and θ_{Vf} are the wage bill and cash flow expressed as a share of value added:

$$\theta_{Lf} \equiv \frac{\omega_{Lf}}{(1 - \sum_i \omega_{if})}, \text{ and } \theta_{Vf} \equiv \frac{\omega_{Vf}}{(1 - \sum_i \omega_{if})}. \tag{7}$$

The first term in equation (6) captures the direct link between a firm's return and its ERP. Intuitively, the return to a firm's specific factor will rise if its ERP rises and fall if the ERPs of other firms rise because rising returns for other firms causes them to bid up the wage. Two important properties of the mapping between ERP and factor prices, which we will use later, are that it is linear and homogeneous of degree 1, which means that factor prices will not change if the ERP does not change.

As we prove in the following proposition, movements in cash flow provide a sufficient statistic for the ERP.

Proposition 3. *The log change in the ERP for a firm ($\hat{p}^e_f$) in a specific factors model is given by*

$$\hat{p}^e_f = \theta_{Vf} \hat{r}_f + \theta_{Lf} \sum_{f'} \frac{L_{f'}}{L} \hat{r}_{f'}. \tag{8}$$

The vectors of log changes in firm output prices and markups are given by $\hat{\mathbf{p}} = \mathbf{A_1} \hat{\boldsymbol{r}}$ and $\hat{\boldsymbol{\mu}} = \mathbf{A_2} \hat{\boldsymbol{r}}$, where the elements of the matrices $\mathbf{A_i}$ are combinations of the factor and input shares. If the share of total expenditures on intermediate inputs is constant, then

$$\widehat{TFPR}_f \equiv \hat{p}_f + \widehat{TFP}_f = \hat{p}^e_f, \tag{9}$$

where $\widehat{TFPR}_f$ is the log change in the firm's revenue total factor productivity. Finally, the vector of changes in quantity TFP ($\widehat{TFP}$) is given by $\widehat{TFP} = \mathbf{A_3} \hat{\boldsymbol{r}}$.

Proof. See Appendix A.3 □

weighted average of the input tariff changes all divided by the share of value added in sales. Despite the popularity of this approach, Ethier (1977) proves that the Corden tariff-based measure of ERP and the Jones price-based measure of ERP cannot be rigorously linked.

Proposition 3 proves that the ERP is simply TFPR. The intuition for this result stems from the fact that cash flow equals the payments to the firm's specific factor, which implies that $\hat{p}^e_f = \theta_{Vf}\hat{r}_f + \theta_{Lf}\hat{w}$. The left-hand side will only be positive if aggregate payments to factors rise, which can only happen if a firm's revenue is growing faster than its costs, i.e., TFPR is rising. The model also yields closed form solutions for how prices and markups move in response to cash-flow movements. For example, if a researcher had access to firm-level intermediate input and labor shares, so that they could compute the $\mathbf{A_i}$ matrices, the model would allow the researcher to compute all price and markup movements in the economy.

2.1.2 Production with Stochastic Prices

Thus far, we have been concerned with a comparative statics exercise in which we assume that the log-changes in cash flows ($\hat{r}_f$) are known with certainty. However, it is straightforward to extend the results to a case in which firms have expectations about future prices but the realized prices may differ. We do this by assuming that prices in the second period equal the prices in the first period plus a random shock drawn from known distributions at the start of the second period, i.e., $\ln p^e_{f2} = \ln p^e_{f1} + \psi_f$, where ψ_f is the shock. Let $\hat{\mathbf{p}}^{\mathbf{e}}$ denote the vector of log changes in ERPs between the two periods, and let $\mathbf{\Psi}$ denote the vector of shocks. Then, the expected changes in log ERPs are $E\left[\hat{\mathbf{p}}^{\mathbf{e}}\right] = E\left[\mathbf{\Psi}\right]$, where $E\left[\cdot\right]$ is the expectation formed in the first period. In addition, we denote the variance of $\mathbf{\Psi}$ as $\sigma^2_{\Psi 0} \equiv \text{Var}\left[\mathbf{\Psi}\right]$, which will capture the degree of uncertainty over whether future ERPs will be higher or lower.

We now consider some policy announcement ($\boldsymbol{\tau}$) that is made at the end of the first period. In our case, we assume that $\boldsymbol{\tau}$ is a vector of tariffs that affects future ERP in two ways. First, the policy can cause expectations about future prices to shift such that the expected changes in log ERPs are now $E\left[\hat{\mathbf{p}}^{\mathbf{e}}|\boldsymbol{\tau}\right] = E\left[\mathbf{\Psi}|\boldsymbol{\tau}\right]$, where $E\left[\cdot|\boldsymbol{\tau}\right]$ is the expectation formed in the first period conditional on the announced tariffs. The expected changes in log ERPs caused by the policy are then $\hat{\mathbf{p}}^{\mathbf{e}}\left(\boldsymbol{\tau}\right) \equiv E\left[\hat{\mathbf{p}}^{\mathbf{e}}|\boldsymbol{\tau}\right] - E\left[\hat{\mathbf{p}}^{\mathbf{e}}\right] = E\left[\mathbf{\Psi}|\boldsymbol{\tau}\right] - E\left[\mathbf{\Psi}\right]$. Second, we allow the policy to shift the variance of $\mathbf{\Psi}$ to $\sigma^2_{\Psi\tau} \equiv \text{Var}\left[\mathbf{\Psi}|\boldsymbol{\tau}\right]$, which captures the idea that a tariff announcement might convey an expectation about what tariffs will be but also uncertainty about whether future tariffs will be higher or lower.

The production structure from the previous section and the stochastic prices presented here can be combined to examine how *expected* changes in cash flow caused by the policy map into *expected* changes in wages, ERP, and other variables. Equation (6) tells us that for any firm f, there is a linear, homogeneous of degree 1 mapping, which we denote by $r_f\left(\cdot\right)$, between the expected change in the rate of return for the specific factor and the vector of expected ERP movements. Thus, we can write the expected log change in cash flow due to the policy as a linear function of the expected log changes in ERPs due to the policy ($\hat{\mathbf{p}}^{\mathbf{e}}\left(\boldsymbol{\tau}\right)$):

$$\hat{r}_f(\boldsymbol{\tau}) \equiv E\left[\hat{r}_f|\boldsymbol{\tau}\right] - E\left[\hat{r}_f\right] = E\left[r_f\left(\hat{\mathbf{p}}^{\mathbf{e}}\right)|\boldsymbol{\tau}\right] - E\left[r_f\left(\hat{\mathbf{p}}^{\mathbf{e}}\right)\right] = r_f\left(\hat{\mathbf{p}}^{\mathbf{e}}\left(\boldsymbol{\tau}\right)\right). \tag{10}$$

Similarly, we can use Proposition 1 and equation (10) to write the policy's impact on expected log change in wages as

$$\hat{w}(\boldsymbol{\tau}) \equiv E\left[\hat{w}|\boldsymbol{\tau}\right] - E\left[\hat{w}\right] = \sum_f \frac{L_f}{L}\left(E\left[\hat{r}_f|\boldsymbol{\tau}\right] - E\left[\hat{r}_f\right]\right) = \sum_f \frac{L_f}{L}\hat{r}_f(\boldsymbol{\tau}). \tag{11}$$

By the same logic, the linear mapping between movements in cash flow and all the variables in Propositions 2 and 3 means that all of our production results will hold in expectation.

2.2 Welfare

In this section, we embed our production framework into the structure of a canonical CCAPM as in Campbell et al. (1997). We consider a setup in which consumers maximize their expected utility over two periods $s \in \{1, 2\}$ while facing uncertainty about what their consumption will be in the second period because future prices, cash flows, and wages depend on unexpected shocks to prices. The L identical consumers (workers) in the economy are each endowed with a unit of labor and initially own an equal share of each firm. We continue to assume the production structure in each period is identical to the one given in the previous sections. We also assume goods are not storable across periods. However, we now assume agents can buy or sell shares in firms at a price of x_{fs} before making their consumption decisions in period one. Each of the A_f shares owned by each agent at the end of the first period entitles its owner to a total payoff of $x_{f2} + d_{f2}$ in the second period, where d_{fs} is the per share dividend paid out by firm f in period s. As before, wages are set to clear the labor market within each period.

We let C_s denote aggregate consumption; P_s the consumer price level in each period s; and TR_s the tariff revenue, which is distributed lump sum. Given initial endowments of each firm's shares A_{f0}, the agent's utility maximization problem can be written as

$$\max_{C_1, C_2, A_f} U \equiv \left(\frac{C_1^{1-\Lambda} - 1}{1-\Lambda} + \zeta E \left[\frac{C_2^{1-\Lambda} - 1}{1-\Lambda} \right] \right),$$

$$\begin{aligned} s.t. \quad C_1 &= \frac{w_1 L + \sum_f r_{f1} V_f + TR_1}{P_1} - \frac{\sum_f x_{f1}(A_f - A_{f0})}{P_1} \\ \text{and} \quad C_2 &= \frac{w_2 L + TR_2}{P_2} + \frac{\sum_f (x_{f2} + d_{f2}) A_f}{P_2}, \end{aligned}$$

where Λ is the coefficient of relative risk aversion and $\zeta < 1$ is the subjective discount factor. If we substitute the constraints into the objective function, the first-order condition with respect to asset purchases (A_f) gives us

$$1 = E\left[\zeta \left(\frac{C_2}{C_1} \right)^{-\Lambda} \frac{P_1}{P_2} \frac{x_{f2} + d_{f2}}{x_{f1}} \right] = E\left[M \frac{P_1}{P_2} \frac{x_{f2} + d_{f2}}{x_{f1}} \right], \tag{12}$$

where $M \equiv \zeta \left(\frac{C_2}{C_1} \right)^{-\Lambda}$ is the SDF and $\frac{(x_{f2}+d_{f2})/P_2}{x_{f1}/P_1}$ is the real rate of return on holding shares of firm f between periods one and two. This equation shows the inverse relationship between the SDF and real rates of return. If consumers are impatient (i.e., when the subjective discount factor (ζ) is low), the SDF (M) will also be low, and real rates of returns must be high to induce agents to hold stocks. Similarly, if expected future consumption (C_2) rises relative to current consumption (C_1), agents will need higher returns in order to induce them to forgo consumption in period one in order to obtain dividend payments in period two.

In equilibrium, the dividends paid out by each firm are determined by its cash flow (i.e., $d_{f2}A_f = r_{f2}V_f$) and the net purchases of each firm's shares is zero ($A_f = A_{f0}$). Also, the equilibrium price of shares in the second period has to satisfy $x_{f2} = 0$ since it is the final period. Substituting these values into the budget constraint, we can express the aggregate income in period s as $I_s = w_s L + \sum_f r_{fs}V_f + TR_s$ and log aggregate consumption as

$$\ln C_s = \ln\left(w_s L + \sum_f r_{fs}V_f + TR_s\right) - \ln P_s. \tag{13}$$

Totally differentiating equation (13) yields

$$\hat{C} = \frac{w_1 L}{I_1}\hat{w} + \sum_f \frac{r_{f1}V_f}{I_1}\hat{r}_f + \frac{TR_1}{I_1}\widehat{TR} - \hat{P}, \tag{14}$$

where hats indicate log changes between periods one and two. We can then exploit the linearity of this expression to write the policy's effect on the expected log change in consumption, $\hat{C}(\boldsymbol{\tau}) \equiv E\left[\hat{C}|\boldsymbol{\tau}\right] - E\left[\hat{C}\right]$, as

$$\hat{C}(\boldsymbol{\tau}) = \frac{w_1 L}{I_1}\sum_f \frac{L_{f1}}{L}\hat{r}_f(\boldsymbol{\tau}) + \sum_f \frac{r_{f1}V_f}{I_1}\hat{r}_f(\boldsymbol{\tau}) + \frac{TR_1}{I_1}\widehat{TR}(\boldsymbol{\tau}) - \hat{P}(\boldsymbol{\tau}), \tag{15}$$

where we make use of the linear relationship between expected changes in wages and cash flow given in equation (11); $\widehat{TR}(\boldsymbol{\tau}) \equiv E\left[\widehat{TR}|\boldsymbol{\tau}\right] - E\left[\widehat{TR}\right]$ is the policy's effect on expected log change in tariff revenue; and $\hat{P}(\boldsymbol{\tau}) \equiv E\left[\hat{P}|\boldsymbol{\tau}\right] - E\left[\hat{P}\right]$ is the policy's effect on expected log change in the price index.

If we assume the price shocks are jointly log-normally distributed and the other variables are linear combinations of prices, consumption movements will also be log-normally distributed.[10] In this case, Rao and Jelvis (2022) present a useful formula for computing a money-metric impact of changes in the means and variances of consumption in terms of "certainty equivalent" ($\tilde{C}$) defined as the amount of guaranteed (i.e., risk-free) consumption one would be willing to forgo in exchange for the stochastic consumption. In particular, they show that the certainty equivalent consumption level is given by

$$\tilde{C} = \exp\left\{E\left[\ln C\right] + \frac{\sigma_C^2}{2}\left(1 - \Lambda\right)\right\},$$

where σ_C^2 is the variance of $\ln C$. Let $\tilde{C}\left(\boldsymbol{\tau}\right)$ and $\tilde{C}\left(0\right)$ denote the certainty equivalents to the stochastic consumption in the second period with and without the policy announcement,

[10] We assume the vector of shocks to log prices is jointly normally distributed so that any linear combination of its elements (i.e., shocks to different prices) is normally distributed. Importantly, this implies that log changes in ERPs and cash flows are each normally distributed because they can be written as linear combinations of log changes to prices (see equations (5) and (6)). Second, we assume that log changes in tariff revenue and the consumer price index are also linear combinations of log changes to prices. For example, these two assumptions are satisfied if tariffs are ad-valorem and if preferences over consumption goods are Cobb-Douglas, respectively. Then, log changes in consumption is normally distributed because equation (15) shows that it is a linear combination of variables that are jointly normally distributed.

respectively. Finally, we measure the welfare impact of the policy announcement with the log change in the certainty equivalent between the two cases given by

$$\ln \tilde{C}\left(\boldsymbol{\tau}\right) - \ln \tilde{C}\left(0\right) = \hat{C}(\boldsymbol{\tau}) + \frac{\sigma^2_{C\tau} - \sigma^2_{C0}}{2}\left(1 - \Lambda\right), \tag{16}$$

where $\sigma^2_{C\tau} \equiv \text{Var}\left[\ln C|\boldsymbol{\tau}\right]$ and $\sigma^2_{C0} \equiv \text{Var}\left[\ln C\right]$ are the variances in log consumption obtained with and without the policy announcement, respectively. Intuitively, in the standard case where we assume the coefficient of relative risk aversion (Λ) exceeds one, a policy's impact on risk-adjusted consumption is rising in the policy's impact on expected consumption and falling in the policy's impact on volatility.

Thus, in order to estimate the impact of a policy on welfare, we need to estimate its impact on expected consumption given in equation (15) and consumption variance given in equation (16). Equation (15) in turn requires us to know how the policy affects expected cash flow, tariff revenues, and inflation.

3 Empirical Framework

This section presents a methodology for using the stock price data to estimate a policy's impact on expected cash-flow and consumption volatility, two of the components we need for the welfare calculation. Stock returns can be written as a function of the present discounted value (PDV) of firm cash flows and discount rates (see Appendix equation (C1)). The theory that we have been building is based on the idea that there are three channels through which policy matters for stock returns. First, a policy change may have a direct effect on the PDV of cash flows by requiring a firm to pay a tariff on inputs or raising its output price. Second, a policy may matter because it moves a macro variable (e.g., exchange rates) that affects the PDV of cash flows. Third, a policy might matter because it affects the SDF.

We estimate the effects on expected cash-flow by first estimating the macro and treatment effects on stock returns and second isolating the cash-flow component of the macro effect on stock returns. Section 3.1 develops a method of estimating the effects of policy announcements on stock returns, which we decompose into those due to impacts on latent macro variables ("the macro effect") and those due to direct impacts on exposed firms ("treatment effect"). Second, in order to identify the expected cash-flow effect, Section 3.2 presents a methodology that builds off Campbell and Vuolteenaho (2004) to extract the cash-flow component of the macro effect of a policy announcement on stock returns. We then sum together the macro and treatment effects on expected cash flow to obtain the total cash flow effect. In the final section, we show how these ingredients can be combined to back out the policy's effects on consumption volatility, which we use for the welfare calculation.

3.1 Effect on Stock Returns

We denote the vector of latent macro variables that matter for stock returns by $\boldsymbol{\delta_t} \equiv (\delta_{1t}, ..., \delta_{Kt})$ and assume that they are affected by macro shocks unrelated to the event $(\boldsymbol{\Phi_t})$ as well as by policy announcements $(\boldsymbol{\tau_t})$ on day t, so we can write $\boldsymbol{\delta_t} = \boldsymbol{\delta}\left(\boldsymbol{\Phi_t}, \boldsymbol{\tau_t}\right)$. We assume that the stock returns are additively log separable into macro and treatment effects:

$$\hat{R}_{ft} = \alpha_f + \hat{R}^M\left(\boldsymbol{\delta}\left(\boldsymbol{\Phi_t}, \boldsymbol{\tau_t}\right), \boldsymbol{\beta_f}\right) + \hat{R}^T\left(\boldsymbol{Z_f}, \boldsymbol{\tau_t}\right) + \nu_{ft}, \tag{17}$$

where α_f is the expected return on firm f's stock in the absence of a policy announcement ($\boldsymbol{\tau_t}$) or other macro shocks ($\boldsymbol{\Phi_t}$); $\boldsymbol{\beta_f}$ is a vector of firm characteristics that matter for how macro variables affect firms; $\boldsymbol{Z_f}$ is another vector of firm characteristics (which may or may not be different from $\boldsymbol{\beta_f}$) that affect how a policy impacts firms directly (e.g., an importer paying a tariff as opposed to having a tariff change some macro variable); and ν_{ft} is a mean-zero error term that captures time-varying, firm-specific shocks that are unrelated to policy announcements.

A key feature of this specification is that it allows a policy announcement to move stock returns through two channels. First, it can have an impact on firms by causing a movement in macro variables ($\boldsymbol{\delta}\left(\boldsymbol{\Phi_t}, \boldsymbol{\tau_t}\right)$). For example, we know from equation (12) that the SDF does not have a firm subscript, so its effect will appear in the $\boldsymbol{\delta}$ and not in the treatment effect. Importantly, this shows that movements in stock returns measured by $\hat{R}^M$ can occur either because of movements in macro factors that are not related to expected cash flow (e.g., discount-rate changes) or because tariffs might cause variables like exchange rates to move, which may have an impact on expected cash flow. Second, movements in $\hat{R}^T$ are net of any discount factor changes and therefore move the firm's stock return through the policy's direct impact on its expected cash flow.

We model the components of $\hat{R}_{ft}$ with a statistical factor model. The use of the factor model is a parsimonious way of modeling an implicit relationship between firm returns and macro variables that have differential effects on firms. Since we do not know the set of macro variables that determine movements in stock returns, we assume that these movements can be described by a set of latent macro variables (δ_{kt}):[11]

$$\hat{R}^M\left(\boldsymbol{\delta}\left(\boldsymbol{\Phi_t}, \boldsymbol{\tau_t}\right), \boldsymbol{\beta_f}\right) = \sum_{k=1}^{K} \beta_{kf}\delta_{kt}, \tag{18}$$

where β_{kf} is our estimate of firm f's stock-price sensitivity to latent variable k (its "loading"). Equation (18) is standard in the asset-pricing literature as it nests many common models. For example, one can see that it nests the CAPM by noting that the CAPM sets $K = 1$ and δ_{1t} equal to the market return.

Given a set of policy announcement events, Ω, we use an event study design to identify the direct effect of tariffs on stock returns. For our application, we define the set of U.S. tariff announcement events as Ω^U, the set of Chinese tariff announcement events as Ω^C, and the combined set of U.S. and Chinese events as $\Omega^{UC} \equiv \Omega^U \bigcup \Omega^C$. Let $D_{jt}^{\mathcal{L}}$ be an indicator variable that is 1 if day t falls within an announcement event window of length $\mathcal{L}$ for event j and zero otherwise. In addition, let $D_t^{\mathcal{L}} \equiv \max_{j\in\Omega^{UC}} D_{jt}^{\mathcal{L}}$ be an indicator variable that is 1 if day t falls within an announcement event window of length $\mathcal{L}$ for at least one event in Ω^{UC} and zero otherwise. Under the assumption that during an event

[11]In order to avoid confusion between the term "factor" as used in statistics and the term "factor" as used in the "specific factors model," we will continue to define a "factor" as a factor of production and refer to the econometric term "factor" as a "latent variable."

window, the movement in a latent variable is determined by the tariff announcement, we write the macro effect of the tariff announcement on the stock returns of firm f as[12]

$$\hat{R}^M\left(\boldsymbol{\delta}\left(\boldsymbol{\tau_t}\right), \boldsymbol{\beta_f}\right) = \sum_{k=1}^{K} \beta_{kf}\delta_{kt}D_t^{\mathcal{L}}. \tag{19}$$

For the treatment effects, we assume that there is a set of treatment variables Z_{fi} ($i \in \{1, ..., N\}$) that specify firm characteristics (relevant only during an event window) that might yield differential returns due to their impact on firm cash flow, e.g., whether a firm imports from China. We further assume that the event's impact on the differential return of a firm f on day t can be written as

$$\hat{R}^T\left(\boldsymbol{Z_f}, \boldsymbol{\tau_t}\right) = \sum_{j\in\Omega^{UC}} \sum_{i=1}^{N} \gamma_{ij}Z_{fi}D_{jt}^{\mathcal{L}}, \tag{20}$$

where γ_{ij} is our estimate of the treatment effect of variable i for each day during event window j.

In order to estimate the macro and treatment effects, we adopt a two-step procedure in which we first estimate a factor model, and then use the residuals from the factor model as the dependent variable in a standard event study. To see how we obtain identification, note that if we substitute equations (18) and (20) into equation (17), we obtain

$$\hat{R}_{ft} = \alpha_f + \sum_{k=1}^{K} \beta_{kf}\delta_{kt} + \epsilon_{ft}, \tag{21}$$

where

$$\epsilon_{ft} \equiv \sum_{j\in\Omega^{UC}} \sum_{i=1}^{N} \gamma_{ij}Z_{fi}D_{jt}^{\mathcal{L}} + \nu_{ft}, \tag{22}$$

$$\nu_{ft} \equiv \xi_t D_{jt}^{\mathcal{L}} + \tilde{\nu}_{ft}, \text{ and } E\left[\nu_{ft}\right] = 0. \tag{23}$$

Here, ξ_t is a parameter to be estimated for each day in event window j, and $\tilde{\nu}_{ft}$ is an error that is mean zero on each day.[13] Once we write the equations this way, one can see that equation (18) is isomorphic to equation (20). In particular, for each i and j combination, we can always define a latent macro variable, $\delta_{\ell t}$, that satisfies $\delta_{\ell t} \equiv \gamma_{ij}D_{jt}^{\mathcal{L}}$. In other words, the treatment effect could be written as a special case of the macro effect in which there is a different ℓ for each i and j combination, and the value of the latent macro variables is given by $\delta_{\ell t} = \gamma_{ij}$ during event window j and zero otherwise. In this case, the coefficient on this latent variable would be given by $\beta_{\ell f} = Z_{fi}$.

[12] As we discuss in Appendix C.2, when we implement this procedure, we first orthogonalize the δ_{kt} with respect to all major macro announcements to mitigate biases due to confounding factors ($\boldsymbol{\Phi_t}$).

[13] The reason we include ξ_t comes from our specification of the moment condition (23). Since we assume that $E\left[\nu_{ft}\right] = 0$, our estimation procedure will impose the moment condition that $\frac{1}{FT}\sum_f \sum_t \hat{\nu}_{ft} = 0$, where F denotes the number of firms, T denotes the number of days, and $\hat{\nu}_{ft}$ is our estimate of ν_{ft} given in equation (23). However, this does not imply that $\frac{1}{F}\sum_f \hat{\nu}_{ft} = 0$, so the value of ξ_t is given by $\xi_t = \frac{1}{F}\sum_f \hat{\nu}_{ft}$. In other words, ξ_t captures the fact that even mean-zero errors need not sum to zero on any given day.

The identification of the number of latent macro variables, and hence what variation appears in the treatment effect, depends on how macro variation is defined. We follow Bai and Ng (2008) and estimate the number of latent macro variables by minimizing the following criteria:

$$IC(K) = \ln(\mathcal{V}(K)) + K\left(\frac{F+T}{FT}\right)\ln(\min\{F, T\}), \tag{24}$$

where F is the number of firms; $\mathcal{V}(K) = (FT)^{-1}\sum_{f=1}^{F}\sum_{t=1}^{T}(\hat{R}_{ft} - \alpha_f - \sum_{k=1}^{K}\beta_{kf}\delta_{kt})^2$ is the average variance of the residuals based on the estimation of equation (21) when K factors are assumed; and T is the number of days.[14] Bai and Ng (2002) show that identification of the number of latent variables (K) assumes that the latent variables (δ_{kt}) matter for stock prices "in general," i.e., we identify any latent variable (δ_{kt}) that has a positive variance in the limit as the sample size approaches infinity:

$$\lim_{T\to\infty}\frac{1}{T}\sum_{t=1}^{T}\left(\delta_{kt} - \delta_k\right)^2 > 0, \tag{25}$$

where $\delta_k \equiv \frac{1}{T}\sum_t \delta_{kt}$. However, under the assumption that the $D^{\mathcal{L}}_{jt}$ are only non-zero during a finite number of event dates, we have $\lim_{T\to\infty}\frac{1}{T}\sum_{t=1}^{T}\left(\gamma_{ij}D^{\mathcal{L}}_{jt} - \gamma_{ij}D^{\mathcal{L}}_{j}\right)^2 = 0$, where $D^{\mathcal{L}}_{j} \equiv \frac{1}{T}\sum_t D^{\mathcal{L}}_{jt} = 0$.[15] The isomorphism implies that we can also write this condition in the notation of the factor model as $\lim_{T\to\infty}\frac{1}{T}\sum_{t=1}^{T}\left(\delta_{\ell t} - \delta_{\ell}\right)^2$, where $\gamma_\ell \equiv \frac{1}{T}\sum_t \gamma_{\ell t} = 0$. Thus, the latent macro variables are those variables that matter in general and the impact of variables that just matter in the event windows will appear in the treatment effect. Since the terms on the right-hand side of equation (22) will appear in the error term of equation (21), we can use the residuals from this equation $(\hat{\epsilon}_{ft})$ as the dependent variable in an event study to identify the remaining parameters.

3.2 Effect on Expected Cash-Flows

The previous section explains how we identify the policy's macro and treatment effects on firms' stock returns, but the presence of a SDF implies that not all of the identified effects on stock returns can be attributed to changes in expected cash-flows. In particular, because the stochastic discount rate is common to all assets and has a positive variance, its effect on stock prices will be captured by the macro effect and not the treatment effect. Hence, although we can safely interpret the identified treatment effects as changes to expected cash flows, we need a method to separate the component of the macro effect

[14] For a given value of K, we estimate the latent macro variables (factors) and their loadings in equation (21) using the method of principal components following Bai and Ng (2002). The distributional assumptions are given in Section 2 of their paper (especially Assumption C). Stated loosely in words, their procedure requires mean-zero errors whose expected absolute values are finite; finite intertemporal and cross-sectional covariances of the errors, and that the expected absolute value of the average elements of the variance-covariance matrix is finite. In contrast to the maximum likelihood estimator, the Bai and Ng (2002) method does not require us to write down a likelihood function.

[15] Having a finite number of tariff events is a sufficient condition to enable us to separate the impact of the latent macro variables from the treatment effect, but it is not a necessary condition. For example, if the frequency of non-zero values of δ_{kt} approaches zero in the limit as $T \to \infty$, it could satisfy the condition.

that reflects movements in expected cash-flows from those arising from movements in the discount rate.

As described in Appendix C.2, we follow Campbell and Vuolteenaho (2004)'s procedure and use a vector autoregression (VAR) to decompose returns into cash-flow and discount-rate components. While they applied the methodology to decompose aggregate market returns, we adapt it to decompose movements in our K latent macro variables. This procedure enables us to obtain an estimate of the share (η_k^{CF}) of each latent macro variable's effect on stock returns (i.e., $\beta_{kf}\delta_{kt}$) that is associated with changes to expected cash flows.[16] Thus, $\eta_k^{CF}\beta_{kf}\delta_{kt}$ gives us the expected cash-flow effect of latent macro variable δ_{kt}.

Armed with these estimates, we can identify the cumulative impact of j events on each firm's expected cash-flows. First, we multiply each δ_{kt} by η_k^{CF} in equation (19) and sum across all days to identify the policy's cumulative macro effect on firm f's expected cash flows as

$$\hat{r}_f^M(\boldsymbol{\tau}) \equiv \sum_t \sum_k \eta_k^{CF} \hat{R}^M(\delta_{kt}, \beta_{kf}) = \sum_t \sum_k \eta_k^{CF} \beta_{kf} \delta_{kt} D_t^{\mathcal{L}}, \tag{26}$$

where we drop the t subscript on $\boldsymbol{\tau}$ when we want to indicate that we are computing the effect of all policy announcements. Similarly, we can use equation (20) and sum across all days to write the policy's cumulative treatment effect on firm f's expected cash-flows as

$$\hat{r}_f^T(\boldsymbol{\tau}) \equiv \sum_t \hat{R}^T(\boldsymbol{Z_f}, \boldsymbol{\tau_t}) = \sum_t \sum_{j \in \Omega^{UC}} \sum_{i=1}^{N} \gamma_{ij} Z_{fi} D_{jt}^{\mathcal{L}}. \tag{27}$$

Thus, the total cumulative change to firm f's expected cash-flows caused by the policy is given by

$$\hat{r}_f(\boldsymbol{\tau}) = \hat{r}_f^M(\boldsymbol{\tau}) + \hat{r}_f^T(\boldsymbol{\tau}), \tag{28}$$

which is a key component for calculating the expected log change in consumption in equation (15). In other words, the expected movement in cash flow due to a policy change is composed of how the event affected expected cash flow through the movement of macro variables and its direct impact on exposed firms.

We can then use equation (11) to compute the expected wage impact of the policy change and Proposition 3 to write down an expression for expected TFPR in terms of expected cash flow movements:

$$\widehat{\text{TFPR}}_f(\boldsymbol{\tau}) \equiv E\left[\widehat{\text{TFPR}}_f | \boldsymbol{\tau}\right] - E\left[\widehat{\text{TFPR}}_f\right] = \theta_{Lf} \sum_f \frac{L_f}{L} \hat{r}_f(\boldsymbol{\tau}) + \theta_{Vf} \hat{r}_f(\boldsymbol{\tau}). \tag{29}$$

[16]Unfortunately, we were unable to estimate the VAR model at a daily frequency because it requires some data that is not available at a daily frequency. We therefore were forced to estimate the VAR model at a monthly frequency instead. As a consequence, we could only estimate the average share (η_k^{CF}) of each latent macro variable's effect on stock returns that is associated with changes to expected cash flows as opposed to a time-varying share (η_{kt}^{CF}).

3.3 Effects on Consumption Volatility

Lastly, we can use these estimates to calculate the change in consumption variance. Appendix C.5 explains the details of the calculation, but the intuition is straightforward. We derive a variant of equation (15) that expresses the expected consumption movement due to expected cash-flow changes at a *daily* frequency ($\hat{C}_t$). In order to estimate the variance in consumption using event windows that start the day before any event j and continue one day afterwards, we first estimate the following regression using all days t in an event window:

$$\hat{C}_t = \sum_{s=-1}^{1} \sum_{j \in \Omega^U} \alpha_s^U I_{jts} + \sum_{s=-1}^{1} \sum_{j \in \Omega^C} \alpha_s^C I_{jts} + \epsilon_t^C,$$

where I_{jts} is an indicator variable that is 1 if day $t = j + s$; α_s^U and α_s^C are estimated parameters that tell us the mean movement of consumption on day s in U.S. and Chinese events, respectively; and ϵ_t^C is an error term that tells us how much our estimate of consumption deviated from the mean level for *that day* within an event window. The variance of this error term is our estimate of how much consumption variance there is for a given day within an event window across events ($\sigma^2_{C\tau}$) in equation (16). We then use an analogous procedure to compute the variance of the consumption in the three days before each event window starts to compute the baseline variance in consumption (σ^2_{C0}).

4 Data

4.1 Data Sources and Variable Construction

Our analysis requires data on stock returns, inflationary expectations, exposure to China, balance sheet items, tariff revenues, and event dates. We provide a summary of variable definitions and sources in Appendix B.1. Our stock return data are from the Center for Research in Security Prices (CRSP) provided by Wharton Research Data Services (WRDS) for every trading day in 2016-2019. When we merge the Compustat data with the CRSP data for a balanced panel of firms that report stock returns ($\hat{R}_{ft}$) on every trading day, we obtain a sample of 2,859 firms that cover all sectors.[17]

We obtain several balance sheet items from Compustat. These include employment, gross sales, profits (operating income after depreciation less interest expenses), cost of tangible fixed property used in the production of revenue less accumulated depreciation, and the net value of intangible assets (e.g., goodwill, blueprints, trademarks, etc.). We adjust the Compustat employment numbers so that they reflect each firm's U.S. employment as explained in Appendix B.6. Based on this procedure, our sample of Compustat firms employs 29.2 million workers domestically or 22.7 percent of the number of people employed in the national employment data provided by the Statistics of U.S. Businesses (SUSB, U.S. Census Bureau).

One problem with Compustat data is that it is not representative of the sample of U.S. firms because it over-samples large firms and does not match the U.S. distribution of industries. These sampling issues can create problems when trying to obtain a mapping

[17]When a company issues multiple classes of stocks, we combine their returns by taking their weighted average, weighted by each stock's share of market capitalization within the firm.

from movements in returns for listed firms and national movements in the variables of interest. We address this problem by assuming that the average returns for Compustat firms in a given employment-size and sector bin is representative for the economy as a whole and then reweight the data using the employment shares of each bin in national data. We describe this procedure in detail in Appendix E.1.

Our measures of inflationary expectations are from Abrahams et al. (2016) using a methodology that isolates the inflation expectations component and updated by Richard Crump. These estimates correct for the fact that the difference in yields on a Treasury security and a Treasury inflation-protected security may not reflect expectations because they may be affected by inflation and liquidity risk. We convert these estimates of inflation into implied movements in the price level over a 10-year period following each announcement. We explain the details of the procedure in Appendix C.4.

We consider three ways in which firms are exposed to China: importing, exporting, and foreign sales (either through exporting or subsidiaries). It is important to capture indirect imports that are ultimately purchased by U.S. firms because many firms do not import directly from China but instead obtain Chinese inputs through their subsidiaries or the U.S. subsidiaries of foreign firms. In order to identify the supply chains, we use DUNS numbers from Dun & Bradstreet to merge importers from Datamyne with a list of firms and their subsidiaries from Capital IQ. We use a firm-name match to link firms, subsidiaries, and their suppliers that are reported in Datamyne, Compustat, Bloomberg, and FactSet and identify which firms are trading with China directly or indirectly through their network of suppliers. After matching firms with identical names in two or more datasets, we manually compared firms with similar names to identify whether they are matches. We define "China Revenue Share" to be the share of a firm's revenues in 2018 (either obtained through sales of subsidiaries or exports) that arise from sales in China as reported in FactSet, and we discuss issues related to the quality of the FactSet data and robustness to alternative measures of revenue shares in Appendix B.5.

The Datamyne data used to identify U.S. firms that import from China or export to China have a number of limitations. First, the product level reported is more aggregated than that in the Harmonized Tariff System 8-digit level at which U.S. tariffs are set. While some of the Datamyne data are at the Harmonized System (HS) 6-digit level, much of it is at the far more aggregated HS2-digit level, making it impossible to know what share of a firm's trade was affected by tariffs. We therefore opt to use a binary exposure measure. Our "China Import" dummy is 1 if the firm or its supply network imported from China in 2017 and zero otherwise. We also construct a "China Export" dummy analogously for exports. Second, the Datamyne data only cover seaborne trade. The U.S. Census data reveal that in 2017, 62 percent of all imports from China and 58 percent of exports to China were conducted by sea. So although we capture over half of the value of U.S.-China trade, the China import and export dummies are likely to miss some U.S. firms that trade with China. On the export side, any exporters that are not reflected in the export dummy are included in the China revenue share variable. To check for missing importers, we also include a robustness check where we replace the importer dummy with a large firm dummy equal to 1 for all firms with more than 1000 employees from Compustat.

These data show that the supply chain information is critical in understanding firms' exposure to international trade. From Table 1, we see that only 10 percent of the firms

Table 1: China Trade Exposure of Listed U.S. Firms

	Mean
Firm imports from China	0.10
Firm or subsidiary imports from China	0.24
Firm, subsidiary, or supplier imports from China	0.29
Firm exports to China	0.02
Firm or subsidiary exports to China	0.04
Firm sells in China via exports or affiliates	0.43
Average share of revenue from Chinese exports or affiliate sales	0.04
Firm exposed to China through imports, exports, or affiliate sales	0.53
Number of Firms: 2,859	

Note: This table reports the means of indicator variables that are 1 if a firm satisfies the listed criterion, as well as the mean of the continuous Chinese revenue share variable.

in our sample import directly from China, and only 2 percent export directly to China. However, if we take subsidiaries into account, these numbers rise to 24 and 4 percent, respectively. When we add imports by all firms in the supply chain, we see that 29 percent of all listed firms in the U.S. import directly or indirectly from China. In the last row of the table, we construct a variable, "Firm Exposed to China" if any firm in the firm's network exported to or imported from China or if the firm had positive revenues from China (possibly from affiliate sales). We see that 53 percent of all firms were exposed to China through one or more of these channels.

In order to compute the change in expected consumption in equation (15), we also need estimates of the announcements' effects on expected tariff revenues. We compute the change in expected tariff revenue ($\widehat{TR}(\boldsymbol{\tau})$) due to the trade-war announcements by using the import demand elasticities estimated in Fajgelbaum et al. (2020) to estimate import quantities (based on Census data) after the levying of the tariffs and multiplying the implied import levels by the amount of the tariff increase. That is, we first construct the counterfactual value of imports that would arise if the only change were the tariffs: $\widehat{\text{Imports}}_{ih,19} = \text{Imports}_{ih,17}(1 - \sigma\hat{\tau}_{ih})\text{Imports}_{h,17}$, where $\text{Imports}_{h,17}$ is the value of imports in 2017 in the HTS10 code h from country i, $\sigma = 2.3$ is the elasticity of import demand estimated in Fajgelbaum et al. (2020), and $\hat{\tau}_{ih} = \ln(1+\tau_{ih,19}) - \ln(1+\tau_{ih,17})$ is the log change in U.S. tariffs due to the policy. We then compute the tariff revenues in 2019 and 2017 as $TR_{19} = \sum_{ih} \widehat{\text{Imports}}_{ih,19}\tau_{ih,19}$ and $TR_{17} = \sum_{ih} \text{Imports}_{ih,17}\tau_{ih,17}$, respectively. Finally, we compute the expected log change in tariff revenues as $\widehat{TR}(\boldsymbol{\tau}) = \ln(TR_{19}) - \ln(TR_{17})$.

4.2 Trade-War Announcements

Over the course of the trade war, the U.S. implemented tariffs in waves, which we plot in Figure 1. The figure shows that the average rate of tariffs on all U.S. imports rose by approximately 4 percentage points as tariffs on a wide range of Chinese imports reached 25 percent by the end of the period.

Figure 1: Average U.S. Tariffs in the 2018-2019 Trade War

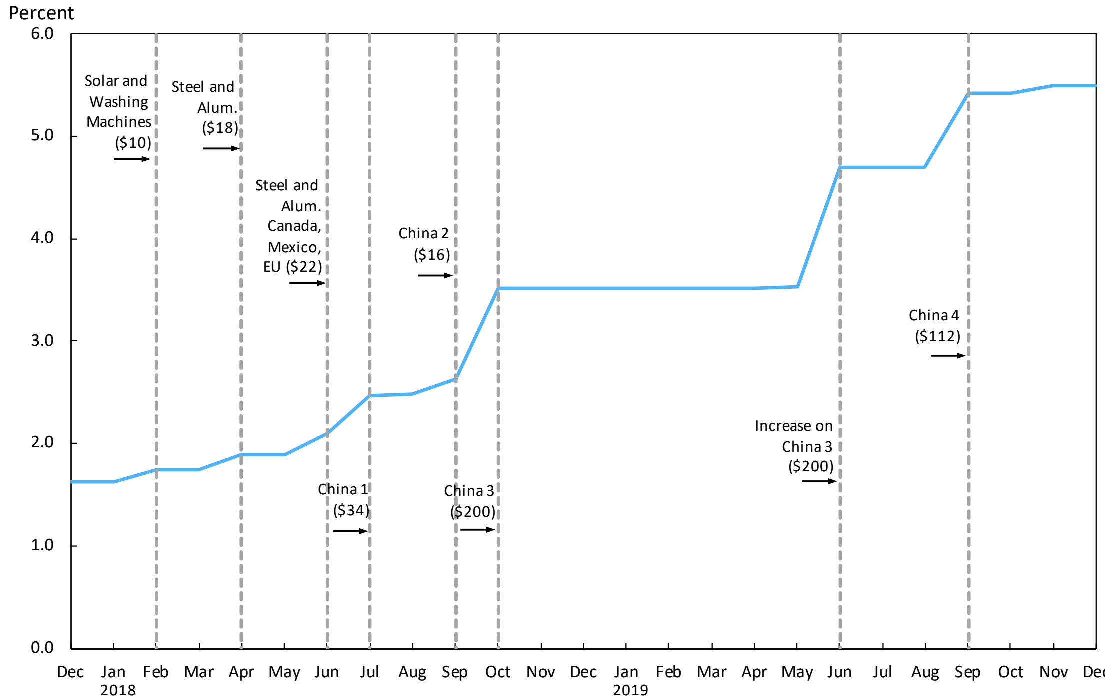

Note: Authors' calculations based on data from the U.S. Census Bureau, U.S. Trade Representative (USTR), and U.S. International Trade Commission. Tariffs on the 10-digit Harmonized Tariff Schedule (HTS) product code by country, weighted by 2017 annual import value. Dashed vertical lines indicate the implementation of new tariffs during 2018-2019; tariffs implemented after the 15th of the month are counted in the subsequent month. Four tranches of tariffs were imposed on China, designated by 1, 2, 3, and 4. Numbers in parentheses correspond to the value of imports covered by the new tariffs in billions.

For each of these new tariffs we identified the earliest announcement date in the media using Factiva and Google search. In addition, we also used the same method to identify the earliest announcement dates for each time that China imposed retaliatory tariffs on U.S. exports. Our method identifies 11 trade-war announcement dates, comprising six U.S. tariff events and five China retaliation events, summarized in Table 2. Our first event is the January 22, 2018 announcement of U.S. tariffs on solar panels and washing machines that were implemented on February 7, 2018 on China and, in this case, more broadly on other countries too. The second event date, the announcement of steel and aluminum tariffs on February 28, 2018, also more broadly applied, was imposed on March 3, 2018. All of the subsequent U.S. tariff events only apply to China. On May 29, 2018 the U.S. announced a 25 percent tariff on $50 billion of Chinese imports. Although this was implemented in two tranches on two separate dates ($34 billion on July 7, 2018 and $16 billion on August 23, 2018), we include this as only one event, since what is important for our purpose is the first time it was announced. All 11 events are listed in Table 2 in date order, with more details and links to the announcement of each event provided in Appendix B.3. Our approach to choosing event dates has the advantage of being comprehensive and objective.

The data reveal that there were large and persistent movements in stock prices and

Table 2: Stock Returns on Event Dates

Event Group	Event Date	R_t (%)	$\sum_{t-1}^{t+1} R_t$ (%)	Description
US	22Jan18	0.75	1.56	U.S. imposes tariffs on solar panels and washing machines
US	28Feb18	-1.07	-3.56	U.S. will impose steel and aluminum tariffs
CHN	22Mar18	-2.57	-4.77	Trade war escalates as China says it will impose tariffs on 128 U.S. exports
US	29May18	-1.00	0.10	White House to impose 25% tariff on $50B worth of Chinese goods
CHN	15Jun18	-0.10	0.08	China announces retaliation against U.S. tariffs on $50B of imports
US	19Jun18	-0.41	-0.28	U.S. announces imposition of tariffs on $200B of Chinese goods
CHN	02Aug18	0.59	0.83	China announces tariffs on $60B of U.S. goods
US	06May19	-0.41	-1.00	U.S. to raise tariffs on $200B of Chinese goods up to 25%
CHN	13May19	-2.58	-1.34	China to raise tariffs on $60B of U.S. goods starting June 1
US	01Aug19	-1.00	-2.90	U.S. will impose a 10% tariff on another $300B of Chinese goods
CHN	23Aug19	-2.64	-1.65	China retaliates with higher tariffs on soy and autos
US+CHN	all	-10.43	-12.94	

Note: This table shows market returns on and around trade-war events. "US" refers to events involving an announcement of U.S. tariffs on China; "CHN" refers to events involving Chinese retaliatory tariffs. R_t is the market return (in our sample of firms) on the day of the announcement. $\sum_{t-1}^{t+1} R_t$ is the cumulative market return over a three-day window beginning on the trading day before the announcement and extending one trading day after. The total three-day return for the U.S. and Chinese events in this table does not exactly equal the value in subsequent tables because we are presenting raw data in this table and double count one day that appears in two event windows.

inflationary expectations following these trade-war announcements. Table 2 presents the stock-market return on each of these event dates. We see that the stock market fell on all of the event dates except one U.S. event date and one Chinese event date, with a total drop of 10.4 percent over all of the events, and 12.9 percent over the three-day windows (beginning the day before the announcement and extending one day after).

We explore the persistence of these stock-market movements in Figure 2, which plots the cumulative log change in average stock prices starting six trading days before each announcement against the number of days before or after each event. The data reveal that in the five trading days before our events, stock-price movements were quite small. Indeed, there is little evidence of anything out of the ordinary happening in the market before the announcements. However, on the announcement days, just as in Table 2, we see that there was a large decline of over 10 percent. Moreover, it is also quite striking how persistent this decline is; the market did not recover even after five trading days. Thus, there is little evidence that markets overreacted and bounced back from their initial negative assessment of the trade war on expected returns.

Next, we also explore the trade-war announcements' impact on other macro variables. We choose three that we think are likely related to changes in trade-policy: changes in the expected price level, exchange rates, and uncertainty (as measured by the VIX). In order to generate the price-change plot, for each day t, we compute the change in expectation on day t of the price level 10 years in the future, where the definitions of these variables are given in Appendix C.4. We then compute the total change (summing across all events) for each day within a 10-day window around each event. We plot the values in Figure 2 (see Appendix B.4 for details and formulas). We see that in the five days leading up to each announcement, the expected price level 10 years later was within about 50 basis points of the level six days before the announcement. However, on the day of the announcements,

Figure 2: Impact of Trade-War Announcements

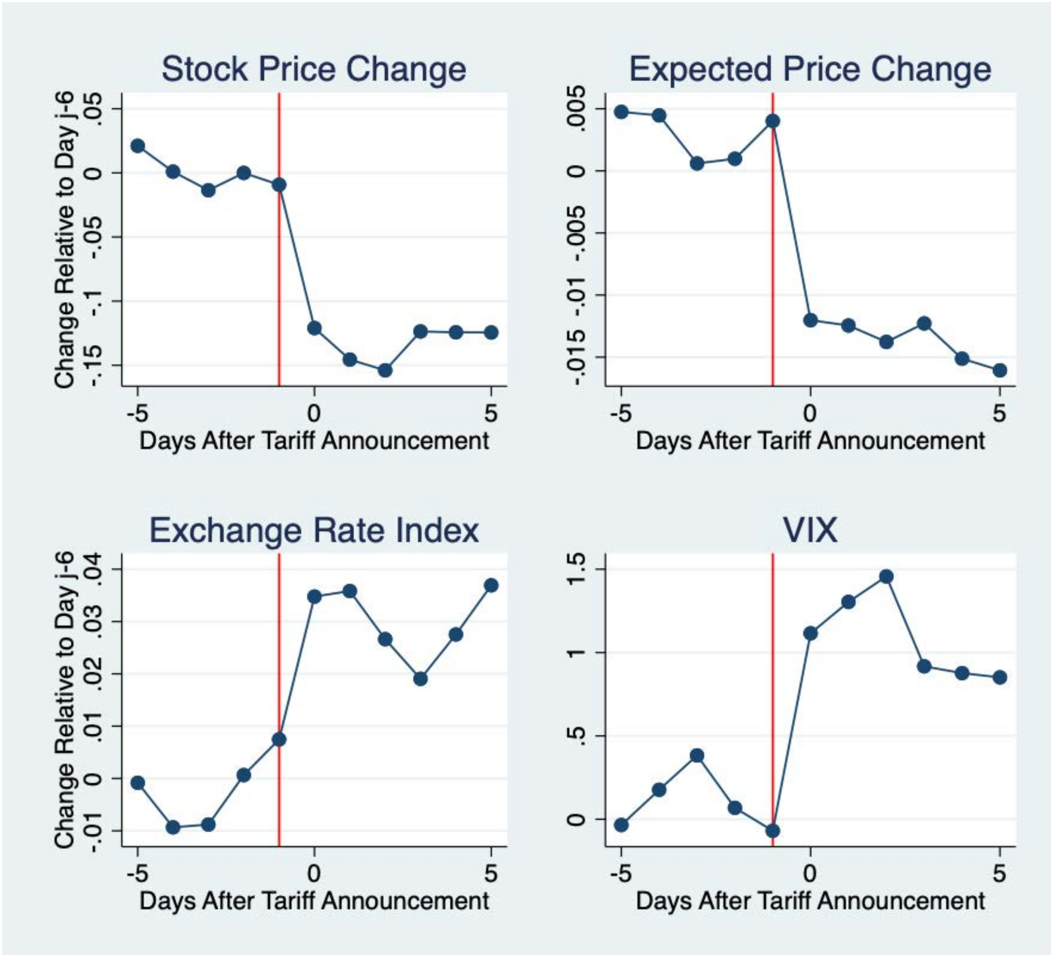

Note: Each chart shows the cumulative log change in a variable starting five trading days before each announcement. Details on this procedure are provided in the text and Appendix B.4. The exchange rate index is the Federal Reserve's Trade Weighted U.S. Dollar Index: Broad, Goods and Services). The exchange rate index is measured in foreign currency per dollar, so higher values correspond to dollar appreciation. Both the VIX and exchange-rate data were downloaded from FRED.

the price level fell sharply, falling close to 150 basis points from the day before. Moreover, there appears to be no recovery in inflationary expectations. If anything, the data suggest inflationary expectations fell as time went on, indicating that just as we saw in the stock-price data, the announcements were associated with a sharp and persistent decline in expected aggregate prices as well.

The lower two panels in Figure 2 present analogous plots for movements in the trade-weighted exchange rate and the VIX. The exchange-rate index is measured in foreign currency per dollar, so higher values correspond to dollar appreciation. As conventional theory would predict, tariff announcements were associated with a 3.3 percent appreciation of the dollar. These trade announcements also caused a 115% increase in the value of VIX, consistent with a rise in uncertainty. There is also no evidence of a speedy mean reversion as these changes persisted for at least 5 trading days after the events. Thus, tariff announcements had significant impacts on a variety of macro variables related to prices and uncertainty.

As we explain in Appendix C.4, we use an event study to measure the impact of tariff announcements on expectations about future aggregate price changes, which we use in our welfare calculations. We report the results of our estimates of the trade war's effect on the expected future price level in Table 3. The estimates indicate that during each day in an announcement window, trade-war announcements were associated with a 0.029 percentage point drop in the expected price level five years later and a 0.038 percentage point drop ten years later. Given that we have 11 events spanning three days each, our results imply that the trade war lowered inflationary expectations so that prices were expected to be 0.96 percentage point lower 5 years later ($= 33 \times -0.029$) and 1.3 percentage points lower 10 years in the future. In columns 2 and 4, we investigate whether it is U.S. or Chinese events that led to the decline in the expected U.S. inflation by interacting the event dummy with a dummy that is 1 if the announcement originated in China. We see that virtually all of the deflationary impact of tariffs comes from U.S. announcements, with Chinese events having no impact on expected U.S. inflation.[18]

Table 3: Impact of Trade-War Announcement on Inflation Expectations

	(1)	(2)	(3)	(4)
	5-year	5-year	10-year	10-year
Event Dummy	-0.029*	-0.076***	-0.038**	-0.092***
	(0.015)	(0.021)	(0.018)	(0.025)
Event Dummy $\times$ China Event Dummy		0.092***		0.105***
		(0.030)		(0.035)
N	989	989	989	989

Note: The dependent variable is the change in inflation expectations 5 years out in columns 1 and 2, and 10 years out in columns 3 and 4. The coefficients reflect the average effect across all event days.

5 Results

In Section 5.1, we present our estimates of the impact of tariff announcements on the expected cash flow of firms ($\hat{r}_f(\boldsymbol{\tau})$) in four steps. First, we follow Bai and Ng (2002) to estimate the factor model in equation (21). That is, the latent macro variables (δ_{kt}), their number (K), and the loadings on them (β_{kf}). Second, we orthogonalize the latent macro with respect to our economic surprise variables to eliminate any movements in the δ_{kt} that could be caused by spurious correlations with data releases (see Appendix C.2 for details). Third, we use the Campbell and Vuolteenaho (2004) procedure to isolate the movement in each orthogonalized latent macro variable that is only due to movements in expected cash flow, which enables us to construct the macro effect ($\hat{r}_f^M(\boldsymbol{\tau})$). Fourth,

[18]It may seem surprising that tariffs can be deflationary, but trade models do not contain a theory of the price level, so they provide little guidance as to whether tariffs should cause aggregate prices to rise or fall. One interpretation of our result is that the output gap widened as a result of the tariffs. As we will see in Table 9, our model predicts that consumption falls by more than productivity, which is consistent with aggregate demand forces driving down prices.

we use our event-study estimates to construct the expected impact on cash flows due to the treatment effect ($\hat{r}_f^T(\boldsymbol{\tau})$). The overall impact on expected cash flow is then the sum of these two components. We describe the results from the first three steps in Section 5.1.1 and the event-study results in Section 5.1.2. We also provide a summary of the estimation procedures for each equation as well as the definitions of the variables and parameters in Appendix C.1 as a quick reference for the reader.

Section 5.2 shows that our estimates of the impact of tariff announcements on the expected cash flow of firms ($\hat{r}_f(\boldsymbol{\tau})$) are plausible in that they can be used to forecast actual movements in profits but not movements in firm-specific factors (which we measure using intangible assets). We also show that expected cash-flow movements can be used to forecast movements in output and employment, as proven in Propositions 1 and 2. This establishes the link between the financial variables and observable movements in real variables. We end by presenting our estimates of how tariff announcements affected expected wages, consumption volatility, welfare, and TFPR in Section 5.3.

5.1 Estimating the Cash-Flow Effect

5.1.1 The Macro Effect

We first present the results of estimating the factor model given in equation (21). We estimate that there are four latent macro variables. As expected, each of these latent macro variables exhibits economically and statistically significant correlations with macro variables like exchange rate movements, changes in expected inflation, and the VIX (see Appendix C.3 for details). This result is consistent with our hypothesis that the forces that move our latent macro variables also move variables that are often central to macro models. While it is not possible to precisely link each latent variable with a particular measurable variable, the correlation between the first factor and the overall market return is 0.84, so it likely captures the sensitivity of firms to the forces that move the overall market (i.e., the forces that would be captured by the CAPM). The other factors also exhibit significant correlations with the macro variables portrayed in Figure 2, which is consistent with the notion that the latent variables are capturing the heterogeneous effects that macro variables can have on firms. The first latent variable accounts for 11.1 percent of the firm-level variance, but additional factors account for much less, with the next three factors accounting for 1.7, 1.5, and 0.9 percent of the variance, respectively. Thus, macro variables explain 15.2 percent of the variance in firm returns over the sample period, and no single potentially omitted macro factor can explain more than 0.9 percent of the variance.

Another way of gauging the success of these latent macro variables in explaining stock returns is to compare it with canonical models used in finance. Our four latent macro variables account for approximately 15 percent of the firm-level variance in stock returns while the CAPM, Fama-French 3-factor model, and Fama-French 5-factor model explain 9, 12, and 13 percent of the variance, respectively. Thus, our setup performs relatively well compared with standard empirical financial models in terms of its ability to match the data.[19]

[19]However, it is also important to acknowledge that this superior fit comes at the cost of not knowing what variables are being captured by the latent macro variables.

With these estimates of latent macro variables in hand, we follow the procedure explained in Appendix C.2 to orthogonalize them to economic surprise variables and isolate the movements in these factors due solely to cash flow. Since our methods and results for this decomposition are largely in line with those of Campbell and Vuolteenaho (2004), we do not discuss them here. We then use equation (26) to construct the macro effect ($\hat{r}^{M}_{f}(\boldsymbol{\tau})$).

5.1.2 Estimating the Treatment Effect

Figure 3: Dispersion in Returns (Three-Day Windows)

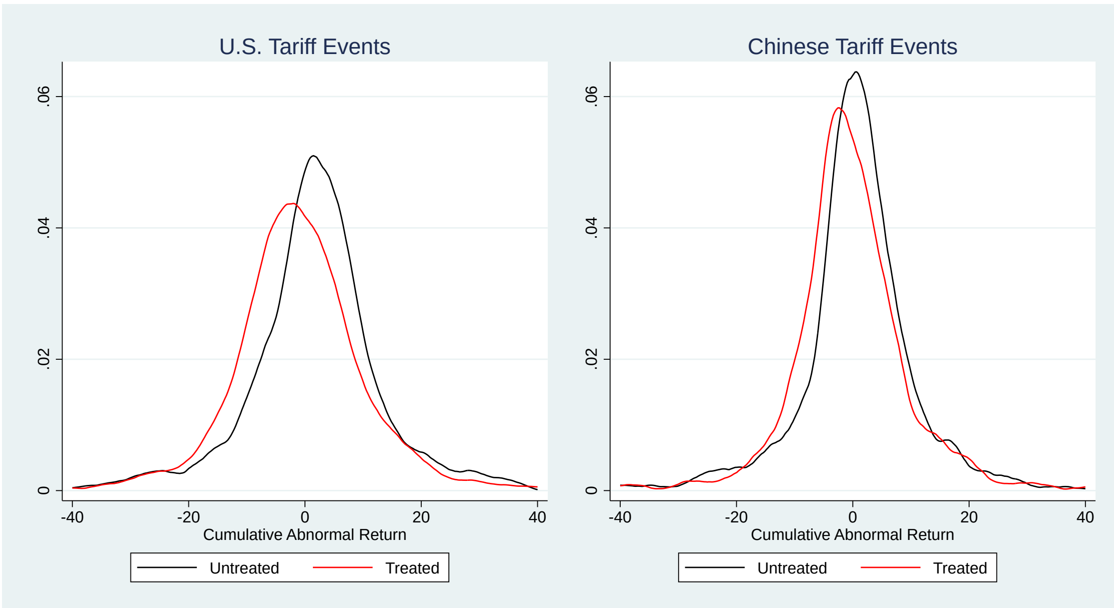

Note: This figure plots the kernel densities of cumulative abnormal returns of firms exposed to China (red) and unexposed (black) during three-day windows around trade-war announcements. Exposed firms are firms that export to, import from, or have positive revenues in China.

Turning to the estimation of treatment effects (i.e., the γ_{ij} in equation 22), Figure 3 shows how the tariff announcements affect firm abnormal returns depending on whether they were exposed or unexposed (see Appendix C.1 for a summary of data definitions). We plot the kernel densities of the cumulative abnormal returns given by the error terms in equation (21) using a baseline three-day window around each event j ($\epsilon_{fj} \equiv 100 \times \sum_{t=j-1}^{j+1} \epsilon_{ft}$) for treated and untreated firms. We define the set of treated firms as firms that import from, export to, or have some positive revenues in China. We see that the distribution of abnormal returns for firms exposed to China during U.S. tariff announcement events is to the left of firms that were not directly exposed. Similarly, we see that announcements of Chinese tariff retaliation produce a similar pattern, with the distribution of abnormal returns for exposed firms lying to the left of the distribution for unexposed firms. These patterns suggest that tariff announcements tend to reduce the abnormal returns of treated firms relatively more.

We identify the relative effects of tariffs on the abnormal returns of exposed firms by estimating equation (22) using a three-day event window ($\mathcal{L} = 3$), where we regress the abnormal return on the firm-exposure variables across all 11 events, allowing separate coefficients for each firm type in each event.[20] Table 4 presents the results for each of the six U.S. tariff events and Table 5 presents the estimated coefficients from the same regression for the five Chinese tariff retaliation events. The estimated coefficients under each event date correspond to the $\hat{\gamma}_{ij}$ in equation (22). Thus, columns 2-7 of both tables are estimated jointly in one regression. The coefficients should be interpreted as the average daily effect of the announcement on the returns of exposed firms during the event window relative to unexposed firms. For example, the coefficient of -0.18 on the China importer dummy in column 3 of Table 4 implies that during the three-day event window around the February 28, 2018 steel and aluminum announcement, firms that imported from China experienced declines in their abnormal returns that were on average 0.18 percentage points lower than other firms *every day* within the window. Thus, their *cumulative* relative decline in stock prices was -0.54 ($= 3 \times 0.18$) percentage points. The numbers in column 1 provide our estimate of the cumulative impact over all U.S. events and all days in the event windows ($3 \sum_j \hat{\gamma}_{ij}$). For example, we can see from the first column of this table that the cumulative impact of the U.S. announcements was to lower the returns of U.S. importers by 1.72 percentage points relative to firms that did not import from China. Similarly, the relative returns of exporters were 2.46 percentage points lower than those of non-exporters, and firm's selling in China saw their returns fall by 0.11 percentage point for every percentage point of revenue they obtained from China. The coefficient on China Revenue Share implies that a firm with the average sales exposure to China corresponding to 4 percent of revenue experienced an abnormal return of -0.4 percentage point across all of the U.S. events.

Table 4 suggests that in general U.S. tariff announcements had negative and significant impacts on the abnormal returns of importers, exporters, and firms selling in China. Although the effects are not precisely measured for every event and measure of exposure, 16 of the 18 event-day coefficients are negative, which indicates that typically firms exposed to China had negative abnormal returns relative to unexposed firms following U.S. tariff announcements. When we sum across all events, the cumulative effect is negative and significant for each type of exposure. Interestingly, U.S. tariff announcements caused negative abnormal returns not only for importing firms but also for firms exporting or selling in China more generally. These negative coefficients on the exporter or sales variables are likely due to three (not mutually exclusive) reasons. The first is that exporters may have anticipated that U.S. tariffs would provoke Chinese retaliatory tariffs, thereby lowering the abnormal return of exporters. Second, market participants may have anticipated that U.S. tariffs would also provoke Chinese retaliatory non-tariff barriers that could lower revenues obtained either by exporting or multinational sales. Third, it is also possible that U.S. tariffs weakened the Chinese economy, which could lower expected profits for U.S. firms selling there.

[20] We exclude the abnormal returns of firms that made an individual announcement between one day before the start of an event window and one day after the event window ended when we estimate equation (22). We do this by obtaining the dates of each firm's individual announcements that are likely to influence its abnormal returns from the Capital IQ Key Developments database. These announcements include those that relate to buybacks (announcements, cancellations), dividends (affirmations, increases, decreases), earnings calls, stock splits, mergers and acquisitions (announcements, cancellations), and follow-on equity offerings.

Table 4: Impact of U.S. Tariff Announcements on Stock Returns

	(1) Cumulative	(2) 22Jan18	(3) 28Feb18	(4) 29May18	(5) 19Jun18	(6) 06May19	(7) 01Aug19
China Importer	-1.72***	-0.00	-0.18***	-0.02	-0.10	-0.12*	-0.14
	(0.56)	(0.07)	(0.07)	(0.06)	(0.07)	(0.07)	(0.09)
China Exporter	-2.46**	0.02	0.03	-0.23***	-0.53***	-0.11	-0.01
	(1.05)	(0.10)	(0.10)	(0.09)	(0.11)	(0.12)	(0.18)
China Revenue Share	-11.32***	-1.10***	-0.20	-0.22	-0.69***	-1.12***	-0.44*
	(1.84)	(0.22)	(0.22)	(0.28)	(0.24)	(0.24)	(0.26)

Note: This table presents the estimated coefficients on the U.S. events obtained from estimating equation (22); the estimated coefficients for the Chinese events are presented in Table 5. Day fixed effects are not reported. The dependent variable ($\hat{\epsilon}_{ft} \times 100$) is the abnormal return obtained from estimating equation (21) with four factors multiplied by 100. China Importer is a dummy that equals 1 if the firm or any of its subsidiaries or suppliers import from China. China Exporter is a dummy that equals 1 if the firm or its subsidiaries export to China. China Revenue Share is the share of the firm's revenue that comes from sales in China reported in percentage points. Column 1 presents the cumulative of the coefficients on each of the U.S. event days. Standard errors are in parentheses. Asterisks correspond to the following levels of significance: *** $p < 0.01$, ** $p < 0.05$, and * $p < 0.1$. The number of observations is 80,674.

Table 5: Impact of Chinese Tariff Announcements on Stock Returns

	(1) Cumulative	(2) 22Mar18	(3) 15Jun18	(4) 02Aug18	(5) 13May19	(6) 23Aug19
China Importer	-0.54	0.09	0.00	0.00	-0.16**	-0.11*
	(0.44)	(0.05)	(0.06)	(0.08)	(0.07)	(0.06)
China Exporter	-1.60**	0.02	-0.08	-0.23*	-0.10	-0.15*
	(0.70)	(0.09)	(0.07)	(0.13)	(0.09)	(0.08)
China Revenue Share	-11.54***	-0.69***	-0.59**	-1.12***	-1.12***	-0.33
	(1.91)	(0.23)	(0.25)	(0.30)	(0.22)	(0.42)

Note: This table presents the estimated coefficients on the Chinese events obtained from estimating equation (22); the estimated coefficients for the U.S. events are presented in Table 4. See the notes to Table 4 for variable definitions and details.

Turning to the Chinese announcements, column 1 of Table 5 shows that in general Chinese retaliation did not have a significant impact on the abnormal returns of firms importing from China, consistent with the idea that while U.S. tariff announcements provoked Chinese retaliation, Chinese retaliation did not provoke new U.S. tariffs. However, Chinese retaliation did produce negative returns for firms exporting to China on five out of the six events and for firms selling in China on all six occasions, though these results are not always statistically significant. Overall, Chinese retaliation announcements led to a significant 1.6 percentage point drop in the abnormal returns of firms exporting to China and another 0.12 percentage point drop for every percentage point increase in a firm's sales in China. The results are economically significant as well. Since Bernard et al. (2007) found that 79 percent of U.S. importers also export, it is worth considering the impact of the trade war on a firm exposed to China through multiple channels. We estimate that a firm that imported from and exported to China and obtained 4 percent of its revenue from sales to China would have had an abnormal return that amounted to -7.2 percent when we sum across all event days. The large magnitude of this result suggests that the trade war had a sizable economic impact on exposed firms.

Table 6: Robustness Tests

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
China Importer			-1.42**	-1.15***	-0.11	0.01	-3.05***	-1.04**
			(0.57)	(0.29)	(0.25)	(1.24)	(0.59)	(0.47)
Large Company	-1.97***	-0.39						
	(0.52)	(0.44)						
China Exporter	-2.62***	-1.76***	-2.50**	-0.93	-0.67	0.23	-3.37***	-1.87**
	(1.01)	(0.67)	(1.06)	(0.67)	(0.41)	(1.71)	(1.09)	(0.78)
China Revenue Share	-11.29***	-11.55***	-10.07***	-6.63***	-5.65***	0.73	-14.11***	-12.31***
	(1.84)	(1.91)	(1.91)	(1.09)	(0.81)	(5.64)	(2.00)	(1.88)
Industry Protected			-0.36					
			(1.28)					
N	80,674	80,674	80,674	29,356	29,356	82,080	80,674	80,674
Event	U.S.	China	U.S.	U.S.	China	Placebo	U.S.	China
Window Size	3	3	3	1	1	3	3	3
Model	4-factor	4-factor	4-factor	4-factor	4-factor	4-factor	CAPM	CAPM

Note: This table presents the results from estimating equation (22) for all U.S.-China tariff events. The dependent variable ($\hat{\epsilon}_{ft} \times 100$) is the abnormal return obtained from estimating equation (21) with four factors multiplied by 100 in columns 1-6, and we use the analogous abnormal return from the CAPM in the last two columns. Large Company dummy equals 1 if a firm had more than 1000 employees in 2017. China Importer is a dummy that equals 1 if the firm or any of its subsidiaries or suppliers import from China. China Exporter is a dummy that equals 1 if the firm or its subsidiaries export to China. China Revenue Share is the share of the firm's revenue that comes from sales in China reported in percentage points. Industry Protected is a dummy that equals 1 if a U.S. tariff was announced in the firm's main 6-digit NAICS industry. Standard errors are in parentheses. The point estimates are the estimated cumulative impact on all event days.

We present a number of robustness tests in Table 6. Each of these specifications is based on estimating equation (22) using a different set of controls. However, in order to save space, we only report the cumulative results, so the columns in Table 6 are comparable with the first column of Tables 4 and 5. We first present results related to measurement error of our trade exposure variables, which could cause attenuation biases leading us to underestimate cash-flow effects. Our coverage of U.S. firms selling in China is likely to be comprehensive because we can identify them either through the export dummy or the China revenue share variable. However, since we can only identify importers if they import goods by sea, our import variable is measured with error and potentially misses some firms that import by air. Bernard et al. (2007) have documented that importers are likely to be large firms. This is also true in our data where we find that 82 percent of importers have a thousand or more employees. In columns 1 and 2, we replace our import dummy with a dummy that is 1 for firms with a thousand or more employees.[21] We find qualitatively similar results, with the coefficient on the large dummy for the U.S. events being negative and significant as in column 1 of Table 4, and smaller and insignificant for the China events as in column 1 of Table 5.

[21]We cannot obtain significant results with both import and large dummies because they are very highly correlated.

Next, we explore whether import tariffs provide protection to firms in that industry by including a dummy equal to 1 if there was an announcement that a new tariff would be levied on imports in the firm's main NAICS 6-digit industry, which would lead to positive abnormal returns for import-competing firms in those sectors. However, the insignificant negative coefficient in column 3 suggests that this is not the case. This finding can be understood by recalling the result in Amiti et al. (2019) showing that U.S. protection drove up domestic output and input prices in treated sectors relative to untreated ones. In particular, ERP could fall if the impact of the tariff on a firm's output price is less than the impact of other tariffs on the pricing of the firm's intermediate input suppliers. Appendix Table D.1, which reports the individual event date coefficients on the Industry Protected variable, highlights this mechanism in our data. It shows that while the only large, *multilateral* application of tariffs—the steel and aluminum tariffs—did cause the abnormal returns of steel and aluminum producers to rise significantly, U.S. tariffs did not help protected industries when they were *only applied bilaterally* against China. Thus, a natural interpretation of this result is that purely bilateral tariffs levied on China raised the prices of Chinese intermediate inputs but failed to afford firms with protection because they still faced competition in their output markets from other foreign suppliers.

Another potential threat to our identification strategy arises from the possibility that our choice of event-window length might be driving the result. For example, while three-day windows allow us to take account of possible information leaks the day before the event or related clarifications after the event, they may also allow for confounding information releases around event days. We deal with this challenge to identification in two ways. First, in columns 4 and 5 of Table 6, we shorten our event windows to one day, so we only consider stock-price movements on the day of the announcement. The coefficient patterns are similar to what we observe using three-day windows. Second, it also might be the case that stock prices bounced back after their initial drops, so in Appendix D.2 we also present results using five-day windows. The results are qualitatively similar to those with the three-day windows and the patterns in Figure 2, which rejects the overshooting hypothesis and suggests that markets became somewhat more pessimistic about the trade war in the days following the announcements. Of course, it may be the case that it took markets more than five trading days to work out that they had made a mistake and overreacted to the tariffs. However, if markets initially overreacted to the initial tariff announcements, this explanation would still leave open the question of why we found in Table 2 that many of the largest stock-market reactions happened in 2019. If markets were irrationally overreacting to the new global situation created by the trade war, it is hard to explain why they would have continued to overreact more than a year after the trade war began.

It is also possible that our results are just due to bad luck—perhaps, we just happened to pick days on which other, non-trade-war related announcements caused the returns of firms exposed to China to fall abnormally. We test the plausibility of this idea by running a placebo test in which we randomly select eleven days out of all trading dates in 2016 to 2019 (excluding our event dates) and re-estimate our event study for each of these randomly chosen events. We repeated this exercise 1,000 times and report the mean coefficients with their associated standard errors in column 6 of Table 6. We find that all coefficients are insignificantly different from zero, which provides another way of rejecting the possibility that our results are just due to chance.

Finally, we explore the role played by the factor model in the last two columns of Table 6. The CAPM setup is commonly used in event studies, but we eschewed its use because we did not want to constrain the way that macro factors affect stock prices. Nevertheless, we can see the role played by our factor model by replacing it with a CAPM framework in which abnormal returns used in the event study are computed based on the CAPM setup. We present these results in columns 7 and 8. The results indicate that using a CAPM setup leads to larger relative effects of trade-war announcements on exposure variables. By including a richer set of macro controls, we tend to obtain smaller estimates of the trade war's effect on treated firms relative to untreated firms.

5.1.3 The Total Cash-Flow Effect

We use the estimates of the macro and treatment effects defined in equations (26) and (27) to decompose movements in expected cash flow (equation (28)). Figure 4 plots the distribution of the log changes in expected cash flow by bins that correspond to the size of the firm and the sector of the economy in which it operates. The market-capitalization weighted average of the macro effect on expected cash flow is -4.4 percent and that of the treatment effect is -2.6 percent, so the total decline in the market that we attribute to the trade war's impact on expected cash flow is 7.0 percent. This is just over half of the actual decline of 12.9 percent that we saw in Table 2, which indicates that a substantial fraction of the aggregate stock movement on these days is associated with changes in expected future cash flows with the remainder due to increased uncertainty and other forces that moved the SDF.

The figure also reveals that there are important differences in the impact of the trade-war on the expected cash-flow of firms of different size and by sector. As expected, the treatment effect is biggest for firms producing goods and for firms employing a large number of workers. This reflects the fact that large firms selling goods are more likely to be buying from or selling in China. Interestingly, we observe a similar firm-size gradient for the macro effect. This result is consistent with the idea that large firms are more likely to be globally engaged, so the trade war is more likely to adversely affect their cash flows. The most striking feature in this figure is the relative magnitude and pervasiveness of the macro effect. The widespread negative effect of the trade war probably reflects the fact that the reduction in expected cash-flows of globally-engaged firms had negative impacts on all firms in their supply chains, regardless of their exposure to China.

One concern with our estimate of the macro effect is that we may be capturing both the effect of the trade-war announcements and some other announcement that coincided with these days. As we noted earlier, the estimates are already purged of the effect of any economic surprises. However, one still may wonder how likely it is that we would have identified a macro effect of this magnitude if we had just randomly picked 11 days between 2016 and 2019. In order to estimate this, we removed each event day from our sample along with the two prior and two subsequent days to create a sample of potential placebo event days in which no trade-war announcement occurred. We then randomly selected 11 placebo event days and their associated event windows, computed the macro effect, and repeated this procedure 1,000 times. We find that in contrast to the 4.4 percent decline in the market due to the macro effect that we estimate for the trade-war

Figure 4: Change in Expected Cash Flow by Firm Size

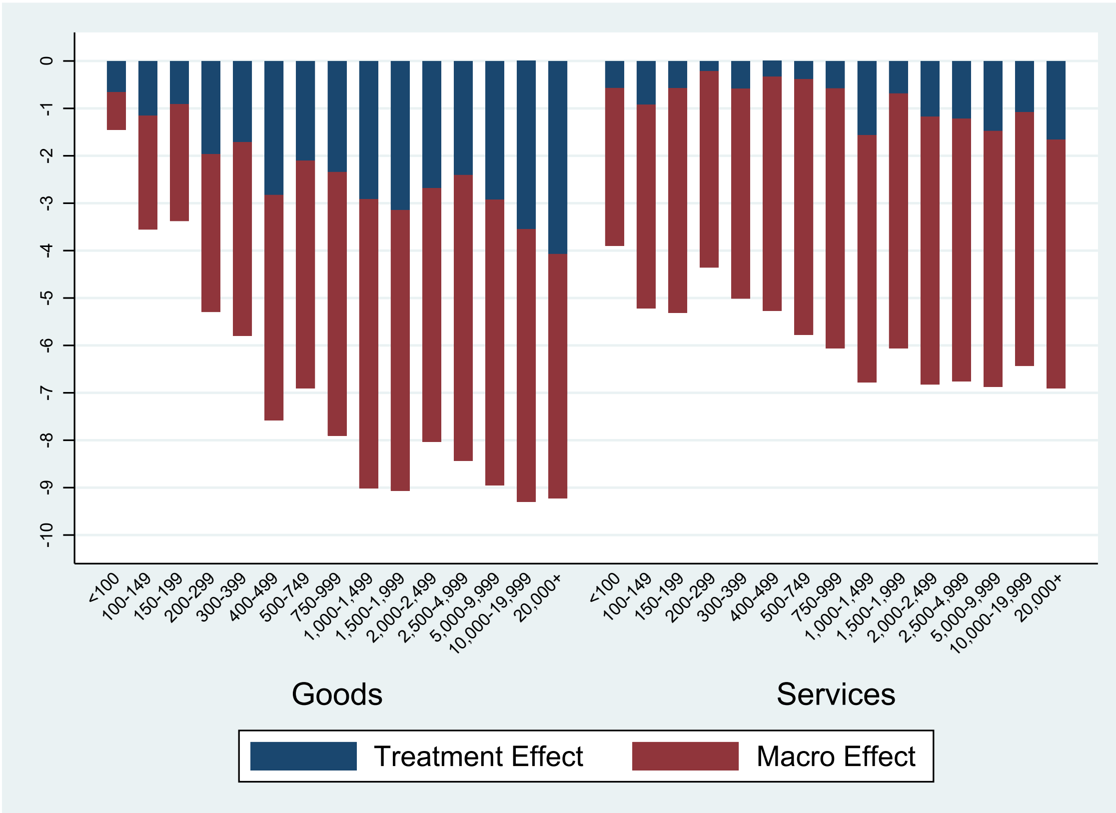

Note: The blue bars correspond to the predicted treatment effects in equation (27) with estimated coefficients from Tables 3 and 4, and the red bars correspond to the predicted macro effect from (26). Both are in percent. These are reported at the sector/firm employment bin level. The weights are adjusted to reflect the economy-wide distribution in 2017, with data on employment distribution by sector-size from the U.S. Census Bureau. Goods sectors are defined as all 2-digit NAICS industries 11, 21-23, and 31-33.

announcements, the average macro effect for the placebo event days was a 0.3 percent *increase* in stock prices. Moreover, out of the 1,000 placebo trials, only 1.1 percent of the draws produced a macro effect of -4.4 percent or less. Thus, we can reject the hypothesis that the macro effect that we identify could have arisen by chance at conventional levels of significance.

5.1.4 The Cash-Flow Impact on Expected TFPR

Equation (29) allows us to obtain a mapping between trade-policy induced movements in expected cash flows and TFPR that we can assess empirically. If we regress movements in firm-level expected TFPR obtained from equation (29) on our China exposure variables, we can interpret the coefficients as an estimate of how exposure to the trade war affected the market's expectation of the change in revenue TFP. Thus, we can see whether we observe the same links between protection and expected TFPR that past studies have identified using primal TFPR. We report the coefficients in Appendix D.4, which show results that are qualitatively similar to earlier micro studies that have used non-financial data to document a positive link between ERP and TFPR.

Figure 5: Expected TFPR by Firm Size

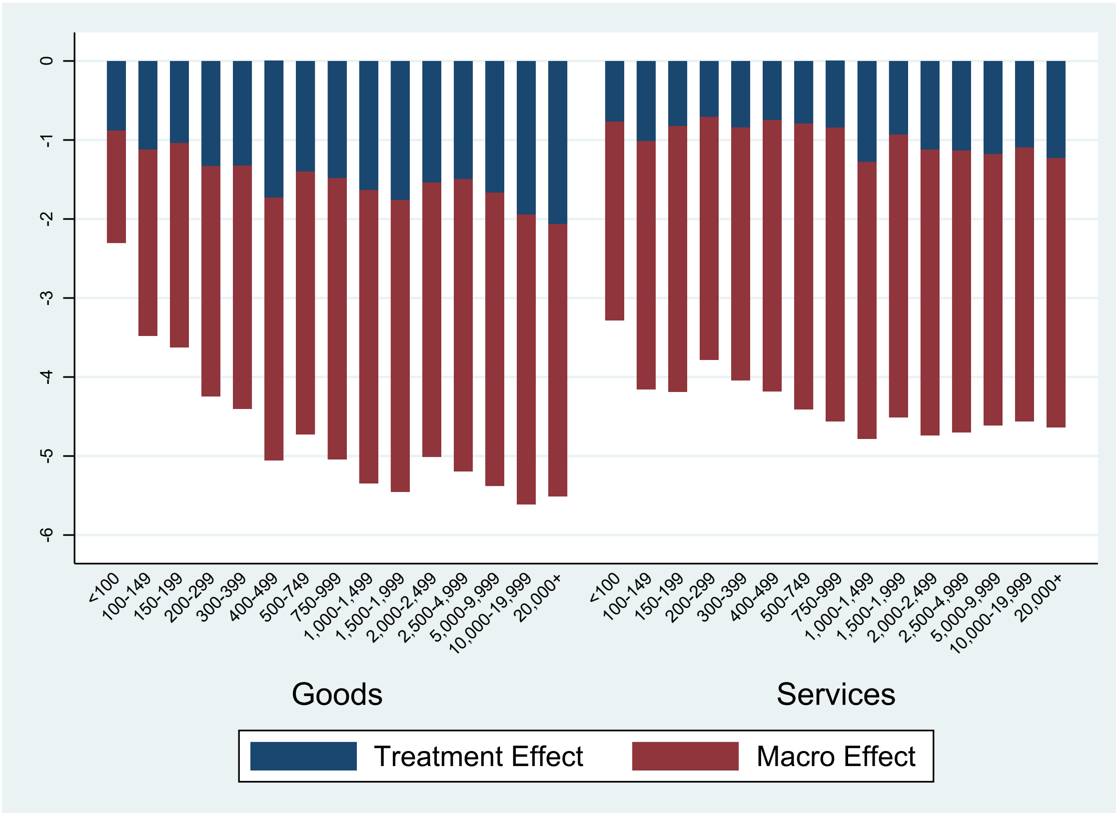

Note: This plot shows estimates of expected TFPR due to movements in the macro effect and the treatment effect by firm size-sector bin. We net out the tariff announcements' effects on aggregate prices in this plot. The precise method is given in appendix equation (E7). Numbers are reported in percent.

However, as we can see in Figure 5, the direct impact of the tariffs on exposed firms is only a fraction of the full impact of the announcements on expected TFPR. Figure 5 shows our estimates of the impact of the tariff announcements on the expected TFPR of firms by firm size. We find that the negative effects of trade-war announcements on expected TFPR increase in magnitude with firm size and then the impact of announcements on expected firm productivity levels off at around -5 percent. Seen through the lens of the specific factors model, the 7.0 percentage point decline in stock prices due to lower cash-flow expectations is associated with a 4.5 percent decline in expected TFPR. Firms employing fewer than a hundred workers experienced expected TFPR declines that were typically 1-3 percentage points less than the declines for large firms. The model suggests that there are two complementary reasons why tariff announcements lowered the expected TFPR of large firms by more. First, large manufacturers of goods are more dependent on trade, so the treatment effect rises in magnitude with firm size. Second, the macro effects—although similar for goods and services firms—also rise in magnitude with firm size. This second story is consistent with the finding in Figure 4 that macro factors drove the share prices of large firms down by relatively more following trade-war announcements.

5.2 The Link Between Expected Cash-Flow and Real Observed Variables

Our theory makes predictions about the link between movements in expected cash flow and future movements in observables that we can use to explore the plausibility of our framework. First, we test whether expected cash-flow movements predict future accounting cash-flow movements and second, whether the amount of each firm-specific factor is unresponsive to movements in expected cash flow. In order to operationalize this test, we need to be clear about what we mean by a firm-specific factor. In particular, a key idea in the finance literature that we use is that there are two broad classes of assets: tangible and intangible. The finance literature models tangible assets as relatively liquid in the sense that they are collateralizable and can be bought and sold. Moreover, it is well-known that listed firms generally finance their tangible asset purchases through borrowing and not through equity.[22] Thus, equity is largely used to finance intangible capital, i.e., trademarks, goodwill, etc.

We can use the fact that intangible assets are difficult to adjust while tangible assets are more adjustable, to both motivate our definition of accounting cash flow and our analysis of which variables should be affected by changes in expected cash flow. Under the assumption that equity finances intangible assets, we should set our "accounting cash flow" to equal operating profits after depreciation and interest payments, where we subtract the last two terms because they largely apply to tangible assets.[23] Similarly, if we assume tangible assets are adjustable (over multiyear horizons), we can model them as just another material input. By contrast, intangible assets are the component of total assets that are harder to adjust because they are not easily bought and sold and end up being largely financed through equity because they are harder to collateralize.

This structure lets us set up our first empirical test of our predictions. We know from Proposition 2 that increases in expected cash flow should increase output, and as long as tangible assets are a "normal" input (i.e., firms use more tangible assets when they produce more output), we should expect a greater cash flow to be associated with increases in investments in tangible assets. However, if the empirical counterpart to a firm-specific asset is its intangible assets, we should expect to see that a rise in the firm's expected cash flow should not affect its stock of intangible assets. We can implement this test by first creating a dummy variable, "Post," that is one for the years after the trade war starts (i.e., 2019, 2020 and 2021) and zero earlier. We then regress each outcome variable on a firm fixed effect, a time fixed effect, and the interaction between Post and our estimate of the expected cash-flow effect due to the tariff announcements ($\hat{r}_f(\boldsymbol{\tau})$). Thus, the coefficient on this interaction term is the difference-in-difference estimate of how much an increase in expected cash flows due to the trade war affected an outcome variable.

We show in Table 7 that our estimates of policy-induced movements in expected cash

[22]For example, Rampini and Viswanathan (2013) "argue that collateral determines the capital structure," and that "that tangibility, when adjusted for leased capital, emerges as a key determinant of leverage and the fraction of firms with low leverage." Similarly, Peters and Taylor (2017) argue that firms do not adjust intangible assets much in response to changing market opportunities. They find that a ten percent increase in investment opportunities (as measured by Tobin's Q) only causes a 3 percent increase in intangible investment—a much smaller number than what they find for tangible assets.

[23]Operating income equals firm revenue less cost of goods sold, sales, general and administrative ex-

Table 7: Relation Between Changes in Expected Cash Flow and Future Observables

	(1) ln(Profit$_{ft}$)	(2) ln(Intangibles$_{ft}$)	(3) ln(L$_{ft}$)	(4) ln(Sales$_{ft}$)	(5) ln(Tangibles$_{ft}$)	(6) ln(Sales/L$_{ft}$)	(7) ln(LP$_{kt}$)	(8) ln(Profit$_{ft}$)
Post$_t$ × Cash-Flow Effect$_f$	1.15**	-0.09	0.56*	1.60***	3.26***			
	(0.50)	(0.75)	(0.31)	(0.41)	(0.50)			
Post$_t$ × TFPR Effect$_f$						2.74***		
						(1.02)		
Post$_t$ × TFPR Effect$_k$							3.81***	
							(1.45)	
Post$_t$ × Discount-Rate Effect$_f$								-1.07
								(1.11)
N	18,009	20,382	24,791	24,184	23,178	24,020	1,219	18,009
Firm FE	yes	yes	yes	yes	yes	yes	no	yes
Industry FE	no	no	no	no	no	no	yes	no
Year FE	yes	yes	yes	yes	yes	yes	yes	yes

Note: The firm-level dependent variables in columns 1 to 6 and 8 come from Compustat. The dependent variable in column 7 is labor productivity at the NAICS 4-digit level from BLS, defined as real output per worker. Data is at annual level for the period 2013 to 2021. Accounting cash flow, "Profit," is operating income after depreciation less interest expenses. L_{ft} is annual employment. Tangible assets are the value of tangible fixed property used in the production of revenue, less accumulated depreciation. Sales are "gross sales." Intangibles are total intangible assets. Post is a dummy variable for years 2019, 2020, and 2021. The Cash-Flow Effect variable is the expected cash flow computed according to equation (28). The TFPR Effect variable in column 6 is the expected TFPR effect computed according to equation (29). In column 7, we aggregate the firm-level data to the NAICS 4-digit industry level using employment weights based on within industry firm size bins from U.S. Census. The Discount-Rate Effect variable in column 8 is computed by using η_k^{DR} instead of η_k^{CF} in equation (26). All dependent variables are multiplied by one hundred.

flow predict the movements in many firm observables in accordance with our theory. We see in column 1 that accounting cash flows in the three years after the start of the trade war are positively associated with our estimate of expected cash flow due to the tariff announcements. In other words, cash-flow expectations based on financial data forecast future accounting cash-flow movements. By contrast, we see in column 2 that the firm's level of intangible assets, our proxy for the amount of each firm's specific factor, are not significantly linked to expected cash-flow movements. Thus, our basic theoretical assumptions that realized cash flow should be positively associated with expected cash flow but investments in intangibles are not associated with expected cash flow is borne out in data.

The following three columns provide support for Propositions 1 and 2 relating future movements in firm employment and sales to movements in expected cash flow due to the trade war. We see in column 3 that as Proposition 1 predicts, there is a positive association between expected cash-flow movements and firm employment (although the coefficient is only significant at the ten-percent level). We obtain stronger support for the relationship between output and expected cash flow derived in Proposition 2: lower expected cash flows due to tariff announcements are significantly associated with lower future sales. Similarly, we also see that lower expected cash flows are also associated with lower levels of tangible assets as the theory would predict for any input that can be adjusted. Although the Compustat data is not well suited for the estimation of TFP, we can use sales per worker as a proxy. We see in column 6 that the expected TFPR variable from Proposition 3 predicts firm revenue productivity. In column 7, we show that if we aggregate our firm-

penses.

level measure of expected TFPR to the 4-digit NAICS industry level, it predicts industry labor productivity. Thus, our three main propositions yield testable predictions about the links between firm-level cash flow and observables that are confirmed in the data.

Finally, we conduct a placebo test to make sure that it is only the cash-flow component of stock returns that is driving the results. As we explain in Appendix C.2, the Campbell and Vuolteenaho (2004) procedure allows us obtain an estimate of the average share of each latent macro variable's effect on stock returns that is associated with changes to expected discount rates $\left(\eta_k^{DR}\right)$. These discount rate movements in stock prices should be unrelated to cash-flow movements and therefore provide the basis for a placebo test to see if this component of stock prices also predicts future accounting cash flow. We implement this test by using η_k^{DR} instead of η_k^{CF} in equation (26) to construct a macro effect composed only of discount-rate movements. If our decomposition approach is successful in separating cash-flow from discount-rate effects, expected discount-rate effects should fail to predict future cash-flow movements. Column 8 confirms that discount-rate effects indeed fail to predict future changes in firms' accounting cash flows, which means that it is only the cash-flow component of stock returns that predicts future accounting cash flows.

5.3 Aggregate Welfare Effects

We now turn to estimating the impact of tariff announcements on the change in expected welfare, which we define as the change in the certainty equivalent in equation (16). We start by computing the impact of the trade war on expected real wages using our estimates of the cash flow movements and adjusting them for the non-representativeness of the Compustat sample (as explained in Appendix E.1). We find that the trade war lowered expected U.S. real wages by 4.3 percent. This economically significant decline arises from two channels. First, the trade war had adverse impacts on productivity operating through the macro and treatment channels, as we showed above, which are expected to depress nominal wages by 5.6 percentage points relative to a benchmark without the trade war. However, this downward pressure on factor prices is also associated with an expected drop in the U.S. price level over a 10-year horizon as we saw in Table 3, and this offsets 1.3 percentage points of the decline.

Next, we check whether the trade war affected welfare by increasing the variance of consumption. As we show in Appendix C.5, by mapping fluctuations in expected cash flows into fluctuations in consumption, we can obtain an estimate of the increase in consumption volatility from any increase in stock market volatility that occurs in the event windows. Our estimates of consumption variance before and after the policy are presented in Table 8, where we set the coefficient of relative risk aversion equal to five because it is what Piazzesi et al. (2007) term the "standard value." Our results show that during the event windows the market's expectation of consumption variance nearly doubled, rising by 0.55×10^{-4} from a base of 0.60×10^{-4} in the three days leading up to each of the event windows.

Table 8: Estimates of Consumption Variance

Description	$\sigma^2_{C\tau}$	σ^2_{C0}	$\Delta\sigma^2_C$
Size-Sector bins	1.16	0.60	0.55
Bins based on size only	1.20	0.76	0.44
Trade war has no impact on small firms	0.96	0.48	0.48
Trade war only matters for listed firms	0.41	0.16	0.25

Note: The reported variances are those of percentage changes in consumption (i.e., variances of log changes in consumption multiplied by 10,000). $\sigma^2_{C\tau}$ reports our estimate of consumption variance after the policy announcements, which is the sum of the consumption variance across all dates within our event windows. σ^2_{C0} reports our estimate of consumption variance before the policy announcements, which is the sum of the consumption variance across all dates within 3-day windows immediately preceding our event windows. $\Delta\sigma^2_C = \sigma^2_{C\tau} - \sigma^2_{C0}$. Details on the calculation of consumption variance are provided in Appendix C.5.

We can assess the plausibility of the estimated values of consumption variance by comparing them to the actual variance in consumption growth, calculated using data on real personal consumption expenditure.[24] First, computing a baseline variance in consumption for non-crisis years using the FRED data, dropping the global financial crisis years 2008 and 2009, we find the variance is almost identical to the baseline consumption variance we obtain from the stock-market data: 0.60×10^{-4} . Thus, our estimate of consumption volatility based on the volatility of the expected cash flows in the days immediately preceding tariff events corresponds to what it actually was for the years excluding the global financial crisis. Second, when we include the global financial crisis years, we find the consumption variance between 2002 and 2017 increases by 1.2×10^{-4}, which is approximately twice the increase in consumption variance between the days in the event windows and those prior that we estimate using stock return data. Thus, our estimate of the implied increase in expected consumption variance is about half of the rise in consumption volatility from its baseline, non-crisis rate to the one we experienced over the fifteen years prior to the trade war. In other words, the stock market data suggests that there was a fairly large, but not historically unprecedented, increase in consumption volatility associated with the trade war.

We report the welfare implications of the trade war in Table 9. We find that expected U.S. welfare fell by 4.9 percentage points as a result of the trade war. The macro effect accounts for about 3.6 percentage points of this drop, with the treatment effect accounting for another 1.3 percentage points, and the rise in the variance of expected consumption accounting for the remaining 0.01 percentage points. In other words, we identify a very small welfare impact arising from the increase in the variance of expected consumption. This result is not because we cannot identify a substantial increase in the variance of expected consumption—the results presented in Table 8 indicate that volatility doubled—but because consumption volatility is sufficiently small that even a doubling of its value has a small welfare effect.

Nevertheless, we explore a number of alternative specifications to see how different assumptions affect our results. In order to see how allowing for sectoral heterogeneity af-

[24] We use a consistent series from FRED (PCEC96). We do not include years following 2017 because they confound the impact of the trade war with the COVID pandemic

Table 9: Aggregate Welfare Effects

Description	Welfare	Consumption			Consumption	TFP	Real Wage
		Total	Macro	Treatment	Volatility		
Size-Sector bins	-4.86	-4.85	-3.56	-1.28	-0.011	-4.53	-4.32
Bins based on size only	-4.96	-4.96	-3.29	-1.67	-0.009	-4.22	-4.13
Trade war has no impact on small firms	-4.51	-4.50	-3.25	-1.25	-0.010	-3.97	-3.71
Trade war only matters for listed firms	-1.48	-1.48	-0.80	-0.67	-0.005	-0.65	-0.43

Note: Welfare is computed according to equation (16). Consumption is computed according to equation (15). The macro and treatment components of consumption are defined in the text following equation (E6). The macro component of consumption only uses the macro cash-flow effect $(\hat{r}_f^M(\boldsymbol{\tau}))$ and the price-level effect $(\hat{P}(\boldsymbol{\tau}))$ from Appendix C.4 to compute consumption. The treatment component of consumption only uses the treatment cash-flow effect $(\hat{r}_f^T(\boldsymbol{\tau}))$ and the tariff revenue effect $(\widehat{TR}(\boldsymbol{\tau}))$ to compute consumption. Consumption volatility equals the second term in equation (16) and is computed using the method described in Section 3.3. TFP is defined in equation (E7) and is computed by subtracting $\hat{P}(\boldsymbol{\tau})$ from the weighted average of firm-level expected TFPR effects in equation (29). The real wage is defined in equation (**??**) and is computed by subtracting $\hat{P}(\boldsymbol{\tau})$ from the expected wage in equation (11).

fects our results, we only use size bins to compute the welfare results instead of size-sector bins and report the results in the second row of Table 9. These results are qualitatively quite similar to our main specification. One of the biggest problems of using the Compustat data to approximate returns in the U.S. economy is that small firms in the Compustat data are likely to have significantly higher profitability than small firms in the U.S. economy. In order to ensure that these differences are not driving our results, we reran our welfare analysis imposing the assumption that the trade war's impact on the returns for firms employing fewer than 100 workers is zero. We report these results in the third row of the table. Not surprisingly, this restriction does lead to a smaller estimate of the impact of the trade war on U.S. welfare, but we still arrive at the conclusion that the trade war lowered U.S. welfare by 4.5 percentage points. In the last row of the table, we consider a conservative assumption: the trade war only affected listed firms. We implement this approach by recomputing the estimated return in each cell after imposing the restriction that the average return for unlisted U.S. firms always equals zero. We still find that the trade war drove down U.S. welfare by 1.5 percent, which is about a third of our baseline estimate. The reason we obtain moderate effects even when we assume that virtually all firms were unaffected by the trade war is that listed firms constitute 22.7 percent of U.S. employment. As a result, when the expected profitability of listed firms declines sharply, as happened during the trade war, this has significant implications for the expected welfare of Americans.

Taken together, these results imply that the 4.9 percentage point decline in welfare arises mainly because the trade-war announcements depressed the expected cash flow of firms through their effects on the latent macro variables, and the negative impacts on the expected cash flows of both treated and untreated firms served to depress wages. Our estimates are high compared to conventional measures, but one possible explanation is that our estimates include forces that are typically ruled out by assumption.[25] The

[25]We examine the plausibility of our estimates by viewing the stock price movements through the lens of a different model: Perla et al. (2021). Their model places much more structure on how trade costs affect

macro effect, for example, is typically omitted in standard analyses. Similarly, while our treatment effect accounts for 1.5 percentage points of the 4.9 percentage point decline in expected welfare, it is still substantially larger than that of Amiti et al. (2019), who found that the welfare loss due the trade war was $79.1 billion, or 0.4 percent of GDP. One likely explanation for the difference is that standard models compute welfare based on the assumption that tariffs will have no impact on within-firm TFP, whereas our estimates allow for changes in expected TFP.

6 Conclusion

There are many challenges facing a researcher trying to assess the welfare impacts of a policy change in general equilibrium. A major contribution of this paper is to provide a new tractable and rigorous methodology to assess the expected impacts of a policy change based on market reactions to policy announcements. The key ingredient is extracting the expected cash flow component of stock price movements due to the policy change. Seen through the lens of a specific factors model, we show this is a sufficient statistic for movements in sales, wages, employment, markups, and both quantity and revenue total factor productivity. Moreover, it allows us to measure welfare effects from shifts in both expected consumption and consumption uncertainty.

Applying our methodology to the U.S.-China trade war, we find a substantially large effect of 4.9 percent in certainty equivalence. This estimate is much higher than those obtained in static trade models, mostly because we allow the policy change to affect macro variables and within-firm productivity. Moreover, our methodology provides a way to identify the macro effect, and thus helps overcome the "missing intercept" problem in studies that try to use estimates from difference-in-difference estimation to infer aggregate effects. We find that the macro effect is four times as large as the treatment effect. In principle, our methodology can be used to analyze any policy change using readily available daily financial data.

References

Abrahams, M., T. Adrian, R. K. Crump, E. Moench, and R. Yu (2016). Decomposing real and nominal yield curves. *Journal of Monetary Economics 84*, 182–200.

Alcalá, F. and A. Ciccone (2004). Trade and productivity. *Quarterly Journal of Economics 119*(2), 613–646.

Amiti, M. and J. Konings (2007). Trade liberalization, intermediate inputs, and productivity: Evidence from indonesia. *American Economic Review 97*(5), 1611–1638.

firms, but it is a useful point of comparison because it can easily be rewritten in a form that allows a researcher who knows how a trade policy affected stock prices to compute the impact on growth and welfare (see Appendix E.2 for details). These estimates imply that the trade-war announcements reduced the average expected firm profit by about 5.6%, which lowers the profit ratio from 1.861 to 1.757 (i.e., $d\overline{\pi}_{rat} = -0.104$). Using the same parameter values as in Perla et al. (2021), these numbers imply a reduction in the economic growth rate of 0.3 percentage points and a welfare decline of 8.1%, which is even larger than our -4.9 percentage point decline.

Amiti, M., S. J. Redding, and D. E. Weinstein (2019). The impact of the 2018 tariffs on prices and welfare. *Journal of Economic Perspectives 33*(4), 187–210.

Amiti, M., S. J. Redding, and D. E. Weinstein (2020). Who's paying for the us tariffs? a longer-term perspective. In *AEA Papers and Proceedings*, Volume 110, pp. 541–46.

Autor, D. H., D. Dorn, and G. H. Hanson (2013, October). The china syndrome: Local labor market effects of import competition in the united states. *American Economic Review 103*(6), 2121–68.

Bai, J. and S. Ng (2002). Determining the number of factors in approximate factor models. *Econometrica 70*(1), 191–221.

Bai, J. and S. Ng (2008). *Large dimensional factor analysis*. Now Publishers Inc.

Baker, S. R., N. Bloom, and S. J. Davis (2016). Measuring economic policy uncertainty. *Quarterly Journal of Economics 131*(4), 1593–1636.

Bernard, A. B., J. B. Jensen, S. J. Redding, and P. K. Schott (2007). Firms in international trade. *Journal of Economic perspectives 21*(3), 105–130.

Bloom, N., M. Draca, and J. Van Reenen (2016). Trade induced technical change? the impact of chinese imports on innovation, it and productivity. *The review of economic studies 83*(1), 87–117.

Brandt, L., J. Van Biesebroeck, L. Wang, and Y. Zhang (2017). WTO accession and performance of chinese manufacturing firms. *American Economic Review 107*(9), 2784–2820.

Brandt, L., J. Van Biesebroeck, L. Wang, and Y. Zhang (2019). WTO accession and performance of chinese manufacturing firms: Corrigendum. *American Economic Review 109*(4), 1616–21.

Breinlich, H. (2014). Heterogeneous firm-level responses to trade liberalization: A test using stock price reactions. *Journal of International Economics 93*(2), 270–285.

Breinlich, H., E. Leromain, D. Novy, T. Sampson, and A. Usman (2018). The economic effects of brexit: Evidence from the stock market. *Fiscal Studies 39*(4), 581–623.

Buera, F. J. and E. Oberfield (2020). The global diffusion of ideas. *Econometrica 88*(1), 83–114.

Caldara, D., M. Iacoviello, P. Molligo, A. Prestipino, and A. Raffo (2019). The economic effects of trade policy uncertainty. *Journal of Monetary Economics forthcoming*.

Caliendo, L., M. Dvorkin, and F. Parro (2019). Trade and labor market dynamics: General equilibrium analysis of the china trade shock. *Econometrica 87*(3), 741–835.

Campbell, J. Y. (1991). A variance decomposition for stock returns. *Economic Journal 101*(405), 157–179. Campbell, jy Meeting of the royal economic soc Mar 27, 1990 Nottingham, england Royal econ soc.

Campbell, J. Y., A. W. Lo, and A. C. MacKinlay (1997). *The Econometrics of Financial Markets*. Princeton University Press.

Campbell, J. Y. and R. J. Shiller (1988). The dividend-price ratio and expectations of future dividends and discount factors. *Review of Financial Studies 1*(3), 195–228. Campbell, John Y. Shiller, Robert J. 1465-7368.

Campbell, J. Y. and T. Vuolteenaho (2004). Bad beta, good beta. *American Economic Review 94*(5), 1249–1275. Campbell, JY Vuolteenaho, T.

Cavallo, A., G. Gopinath, B. Neiman, and J. Tang (2021). Tariff pass-through at the border and at the store: Evidence from us trade policy. *American Economic Review: Insights 3*(1), 19–34.

Chen, L. and X. L. Zhao (2009). Return decomposition. *Review of Financial Studies 22*(12), 5213–5249. Chen, Long Zhao, Xinlei.

Corden, W. M. (1966). The structure of a tariff system and the effective protective rate. *Journal of Political Economy 74*(3), 221–237.

Crowley, M. A., N. Meng, and H. Song (2019). Policy shocks and stock market returns: Evidence from chinese solar panels. *Journal of the Japanese and International Economies 51*, 148–169.

Dix-Carneiro, R. and B. K. Kovak (2017). Trade liberalization and regional dynamics. *American Economic Review 107*(10), 2908–46.

Ethier, W. (1977). The theory of effective protection in general equilibrium: Effective-rate analogues of nominal rates. *Canadian Journal of Economics*, 233–245.

Fajgelbaum, P. D., P. K. Goldberg, P. J. Kennedy, and A. K. Khandelwal (2020). The return to protectionism. *Quarterly Journal of Economics 135*(1), 1–55.

Feyrer, J. (2019). Trade and income-exploiting time series in geography. *American Economic Journal: Applied Economics 11*(4), 1–35.

Fisman, R., Y. Hamao, and Y. Wang (2014). Nationalism and economic exchange: Evidence from shocks to Sino-Japanese relations. *The Review of Financial Studies 27*(9), 2626–2660.

Flaaen, A., A. HortaÃ§su, and F. Tintelnot (2020, July). The production relocation and price effects of us trade policy: The case of washing machines. *American Economic Review 110*(7), 2103–27.

Frankel, J. A. and D. H. Romer (1999). Does trade cause growth? *American economic review 89*(3), 379–399.

Greenland, A., M. Ion, J. Lopresti, and P. Schott (2020). Using equity market reactions to infer exposure to trade liberalization. *Available at SSRN*.

Grossman, G. M. and J. A. Levinsohn (1989). Import competition and the stock-market return to capital. *American Economic Review 79*(5), 1065–1087.

Handley, K. and N. Limão (2017). Policy uncertainty, trade, and welfare: Theory and evidence for China and the United States. *American Economic Review 107*(9), 2731–2783.

Hartigan, J. C., P. R. Perry, and S. Kamma (1986). The value of administered protection: a capital market approach. *The Review of Economics and Statistics*, 610–617.

Hsieh, C.-T. and P. J. Klenow (2009). Misallocation and manufacturing tfp in china and india. *Quarterly Journal of Economics 124*(4), 1403–1448.

Huang, Y., C. Lin, S. Liu, and H. Tang (2019). Trade linkages and firm value: Evidence from the 2018 US-China trade war. *CEPR Discussion Paper DP14173*.

Jones, R. W. (1975). Income distribution and effective protection in a multicommodity trade model. *Journal of Economic Theory 11*(1), 1–15.

Knoblach, M. and F. Stöckl (2020). What determines the elasticity of substitution between capital and labor? a literature review. *Journal of Economic Surveys 34*(4), 847–875.

Kovak, B. K. (2013). Regional effects of trade reform: What is the correct measure of liberalization? *American Economic Review 103*(5), 1960–76.

Lewis, D. J., C. Makridis, and K. Mertens (2019). Do monetary policy announcements shift household expectations? *FRB of New York Staff Report* (897).

Melitz, M. J. (2003). The impact of trade on intra-industry reallocations and aggregate industry productivity. *econometrica 71*(6), 1695–1725.

Moser, C. and A. K. Rose (2014). Who benefits from regional trade agreements? the view from the stock market. *European Economic Review 68*, 31–47.

Perla, J., C. Tonetti, and M. E. Waugh (2021). Equilibrium technology diffusion, trade, and growth. *American Economic Review 111*(1), 73–128.

Peters, R. H. and L. A. Taylor (2017). Intangible capital and the investment-q relation. *Journal of Financial Economics 123*, 251–272.

Piazzesi, M., M. Schneider, and S. Tuzel (2007). Housing, consumption and asset pricing. *Journal of Financial Economics 83*(3), 531–569.

Pierce, J. R. and P. K. Schott (2016). The surprisingly swift decline of US manufacturing employment. *American Economic Review 106*(7), 1632–1662.

Rampini, A. A. and S. Viswanathan (2013). Collateral and capital structure. *Journal of Financial Economics 109*(2), 466–492.

Rao, A. and T. Jelvis (2022). *Foundations of Reinforcement Learning with Applications in Finance*. Taylor and Francis, forthcoming.

Sampson, T. (2016). Dynamic selection: an idea flows theory of entry, trade, and growth. *Quarterly Journal of Economics 131*(1), 315–380.

Topalova, P. (2010). Factor immobility and regional impacts of trade liberalization: Evidence on poverty from india. *American Economic Journal: Applied Economics 2*(4), 1–41.

Topalova, P. and A. Khandelwal (2011). Trade liberalization and firm productivity: The case of india. *Review of economics and statistics 93*(3), 995–1009.