Uncertainty and Trade Agreements*

Nuno Limão   Giovanni Maggi

August 2014

Abstract

We explore conditions under which trade agreements can provide gains by reducing trade-policy uncertainty. Given the degree of income risk aversion, this is more likely when economies are more open, export supply elasticities are lower and economies more specialized. Governments have stronger incentives to sign trade agreements when the trading environment is more uncertain. As exogenous trade costs decline, the gains from reducing tariff uncertainty become more important relative to reducing average tariff levels. We also develop a simple "sufficient statistic" approach to quantify the gains from managing trade-policy uncertainty, and examine the impact of ex-ante investments on such gains.

---

*Limão: Economics Department, University of Maryland, College Park, MD 20817; NBER; email: limao@econ.umd.edu. Maggi: 30 Hillhouse Avenue, Yale University, New Haven, CT 0651; FGV/EPGE and NBER; email: giovanni.maggi@yale.edu. We thank the reviewers for useful suggestions. We also thank Kyle Bagwell, Mostafa Beshkar, Paola Conconi, Swati Dhingra, Avinash Dixit, Robert Staiger for helpful comments as well as participants in the 2012 AEA meetings, the 2012 Venice CAGE/CEP Workshop, the 2012 Princeton IES Summer Workshop, the Australasian Econometric Society and seminars at Stanford, Wisconsin, Oxford, LSE, Warwick, Nottingham and USP.

Policy practitioners often argue that a central benefit of trade agreements (TAs) is to reduce trade policy uncertainty. Indeed, the WTO and many other TAs explicitly state that one of their goals is to increase the predictability of the trade policy environment.[1] But in spite of the importance that policy makers and international institutions attribute to the notion of an uncertainty-reducing role of TAs, we know little about its theoretical underpinnings. A large body of theory has explored the possible roles of TAs as means to correct international policy externalities (e.g. Grossman and Helpman, 1995, Bagwell and Staiger, 1999, and Ossa, 2011) and to allow governments to commit vis-a-vis domestic actors (e.g. Maggi and Rodriguez-Clare, 1998, and Limão and Tovar-Rodriguez, 2011). But this research focuses on the role of TAs in managing the *level* of trade barriers, not their uncertainty.

The main objective of our paper is to explore the conditions under which there is an uncertainty-reducing motive for a TA, and examine the potential gains that a TA can provide by regulating trade-policy uncertainty, above and beyond the more standard gains from reducing the levels of trade barriers.

We focus on a scenario without frictions in contracting between governments, so that the TA is a complete contingent contract. Since we are focusing on the motives and potential gains from a TA, rather than its design, focusing on a setting without transaction costs is a natural first step. To isolate the uncertainty-managing motive for a TA we focus on a thought experiment: the optimal "mean preserving agreement," that is the optimal TA among all agreements that keep the average trade barrier at the same level as in the noncooperative equilibrium. If this agreement leads to a policy distribution that is different from the noncooperative one, we say that there is an "uncertainty-managing motive" (or simply an "uncertainty motive") for a TA, and if it reduces policy uncertainty relative to the noncooperative equilibrium we say that there is an "uncertainty-reducing motive."[2]

[1] For example, the WTO's web site states that "Just as important as freer trade – perhaps more important – are other principles of the WTO system. For example: non-discrimination, and making sure the conditions for trade are stable, predictable and transparent." Several preferential trade agreements such as some of those entered into by the United States, the European Union and by developing countries often claim that they aim to 'reduce distortions to trade' *and* 'ensure a predictable environment for business planning and investment'.

[2] We also consider an alternative thought experiment, which focuses on the tariff schedule that a government would unilaterally choose if it were constrained to deliver the same mean as the optimal agreement. If such "mean-preserving unilateral" choice exhibits more uncertainty than the optimal trade agreement, we say that there is an uncertainty-reducing motive. In section II we discuss the similarities and differences between the results under the two thought experiments, and the reasons

Our first step is to examine a simple framework in which government objectives are specified in reduced form as functions of a trade policy and an underlying shock. Starting from a reduced-form framework with relatively little structure is useful for several reasons. First, the framework delivers general formulas for the direction of the uncertainty motive and the gains from regulating policy uncertainty, which admit intuitive interpretations and make the logic of our results quite transparent. Second, these formulas can be readily applied to a specific trade model to examine how the direction of the uncertainty motive and the associated gains depend on the underlying economic environment. And third, the framework can in principle also apply to other types of international agreements, such as environmental or investment agreements.

Initially we focus on a setting where only one country (Home) chooses a trade barrier, which exerts a negative externality on a policy-passive country (Foreign); later we extend the model to allow for two policy-active countries. The noncooperative level of the trade barrier is increasing in the underlying shock. We identify two key effects that determine whether there is an uncertainty motive for a TA, and if so, in what direction it goes. The first one is what we label the *policy-risk preference* effect, determined by the concavity/convexity of Foreign's payoff with respect to Home's policy: when the Foreign country is policy-risk averse, this effect works in favor of an uncertainty-reducing motive. Intuition might suggest that this effect is all that matters for determining whether there is "too much" or "too little" risk in the noncooperative policy. And indeed this is the case when the shock affects the Foreign country only through Home's policy (a "political economy" shock). However, when the shock affects the Foreign country also in a direct manner (an "economic" shock), there is an additional effect that we label the *externality-shifting* effect. If a higher level of the shock strengthens the marginal international policy externality holding the policy level constant, this effect works in favor of the uncertainty-reducing motive for a TA, otherwise it works against it, and the uncertainty motive need not go in the same direction as the Foreign country's policy-risk preference.

Our next step is to apply these general conditions and formulas to a standard competitive trade model with two countries and two sectors. Focusing on a perfectly competitive setting – rather than one that emphasizes the role of firms – seems like a natural choice, given that this is a first step in exploring the uncertainty motive

why we focus the analysis on the mean-preserving-agreement thought experiment.

for TAs.[3] For simplicity we assume that the Foreign country is small relative to the Home country and maximizes welfare. We allow individuals to be risk averse. In the basic model we consider shocks of the political-economy type, and later extend the model to allow for more general shocks.

It is natural to start with the benchmark case of risk neutral individuals. In this case we find that there tends to be an uncertainty-*increasing* motive for a TA. The reason is that, given the political-economy nature of the shock, all that matters for the uncertainty motive is the Foreign country's policy-risk preference (as mentioned above), and in the presence of income-risk neutrality the Foreign country tends to be policy-risk *loving*. This is due to the convexity of the indirect utility function and of the revenue function in prices, reflecting the ability of firms and consumers to make decisions after observing prices. Interestingly, then, the standard trade model with risk neutrality seems at odds with the often-heard informal argument that TAs can provide gains by reducing trade-policy uncertainty.

When we allow individuals to be risk averse, we find that the uncertainty-managing motive for a TA is determined by a trade-off between risk aversion and flexibility: on the one hand the degree of risk aversion, in interaction with the degree of openness, pushes toward an uncertainty-reducing motive; on the other hand the degree of flexibility of the economy, which in turn is determined by the export supply elasticity and the degree of specialization, pushes toward an uncertainty-increasing motive. We note that, empirically, lower-income countries tend to have lower export supply elasticities and a lower degree of diversification, thus at a broad level our model suggests that the uncertainty-reducing motive for a TA should be more important for lower-income countries than for higher-income countries.

The uncertainty motive for a TA is affected in interesting ways by changes in exogenous trade costs (e.g. transport costs). We show that, if individuals are sufficiently risk averse, as the trade cost declines from its prohibitive level, initially there is an uncertainty-increasing motive for a TA, but this turns into an uncertainty-reducing motive as the trade cost continues to fall. Thus the model suggests that uncertainty-reducing motives for TAs are increasingly likely to emerge as trade costs decline over time, and are more likely to be present for countries within a region.

[3]It is also important to note that, as we discuss in footnote 25, a key effect highlighted by our model would operate also in models with imperfect competition and sunk investments, namely the feature that in the presence of risk neutrality the model tends to predict an uncertainty-*increasing* motive for a TA.

Next we examine the gains from regulating trade-policy uncertainty relative to the gains from regulating the tariff mean. We isolate the latter by focusing on "uncertainty preserving" agreements,[4] while we capture the former (as explained above) by focusing on mean-preserving agreements. We find that, if individuals are sufficiently risk averse, the relative gains from regulating trade-policy uncertainty are non-monotonic in the trade cost: as the trade cost falls, these relative gains initially decrease, but eventually they grow in magnitude. Another implication of the model concerns the impact of the underlying degree of uncertainty, as captured by the variance of the political-economy shock: an increase in this variance leads to larger overall gains from a TA, thus suggesting that governments should have a higher propensity to sign TAs when the trading environment is more uncertain.

Next we extend the model to allow for more general economic shocks. As mentioned above, economic shocks may amplify or reduce the impact of Home's protection on Foreign, thereby introducing a policy-externality-shifting effect, in addition to the policy-risk-preference effect. The externality-shifting effect operates through two possible channels: first, to the extent that the shock affects domestic economic conditions in the Home country, it will affect the Foreign country through the terms-of-trade; and second, to the extent that the shock affects domestic economic conditions in the Foreign country, it will have a further impact on this country. We discuss conditions under which the externality-shifting effect strengthens the uncertainty-reducing motive for a TA.

Our model suggests a simple "sufficient statistic" approach to empirically determine the direction of the uncertainty motive for a TA and to quantify the relative gains from regulating trade-policy uncertainty. We start from the observation that the international externality exerted by Home's tariff is given by an adjusted measure of Foreign's openness (where the adjustment factor is related to the marginal utility of income), and show that there is an uncertainty-reducing motive for a TA if and only if the adjusted openness co-varies with the non-cooperative tariff level as a result of the underlying shocks. Furthermore, the covariance between the tariff and the adjusted openness can be used, in conjunction with the tariff mean, to approximate the relative gains from regulating trade-policy uncertainty.

[4] Specifically, an uncertainty-preserving agreement is an agreement that shifts the tariff schedule in a way that changes the mean but preserves all the higher central moments (variance, skewness, kurtosis, etc.).

We then illustrate how our "sufficient statistic" approach can be taken to the data. We focus on the trading relationship between a small and a large country, Cuba and the US respectively, during a period of non-cooperative trade policies, namely before the 1934 agreement between these two countries. We find a positive correlation between US tariffs and Cuban adjusted openness when calculated at reasonable levels of risk aversion, which suggests there was an uncertainty-reducing motive for a TA between these two countries, and we find that the relative gains from reducing trade-policy uncertainty were significant. Our model is extremely stylized, so this exercise should be interpreted with caution, but it suggests that the model can be taken to the data in a meaningful way, and it points to a potential direction for future research: developing richer versions of the model and taking them to richer datasets.

In our baseline model, factors can be allocated only after uncertainty is resolved. In section VI we extend the model to allow for ex-ante investments. We show that the condition determining the direction of the uncertainty motive for a TA in the presence of ex-ante investments is analogous to the one derived in the static model, provided the market allocation of capital is efficient given Home's trade policy. Even though the TA can change the allocation of capital, this change has no first-order welfare effect in the Foreign country, due to the initial efficiency of the factor allocation. We interpret this result as suggesting that there is no separate uncertainty motive associated with ex-ante investment. Next, we examine the direction in which a TA affects investment and trade *via* changes in policy uncertainty. We show that, if risk aversion is sufficiently strong, the support of the shock is sufficiently small and the export supply elasticity does not increase too steeply with the price, then (i) there is an uncertainty-reducing motive for a TA and (ii) a reduction in policy uncertainty leads to more investment in the export sector and a higher expected trade volume. Overall, our analysis of ex-ante investments suggests an important caveat to the statements made by the WTO and other TAs that an important goal is to reduce policy uncertainty in order to increase investment in export markets: even though a reduction in policy uncertainty does (under some conditions) have this effect, this in itself is not sufficient to ensure a first-order increase in welfare.

Finally, we extend the analysis to allow for two (symmetric) policy-active countries. The general condition that determines the direction of the uncertainty motive for a TA in this case still includes the policy-risk-preference and externality-shifting effects, but now there is an additional effect, which works in favor of an uncertainty-reducing motive if tariffs are strategic substitutes, and against it if they are strategic complements.

Before proceeding, it is natural to ask: empirically, is there significant uncertainty in trade policies? At a broad level one can distinguish between two types of trade-policy uncertainty. A first type is the risk of major protectionist events, such as trade wars, which are very infrequent and thus not typically observed in a specific period of time. A second type is the presence of frequent, small-magnitude changes in trade policy (we will refer to this as trade policy "volatility"), which can be more easily observed in the data.[5]

Regarding the first type of trade-policy uncertainty, it is difficult to measure it empirically, but there seems to be a definite belief among policy makers and businesses that this risk is real, and such belief seems reflected in the stated goals of TAs mentioned above. A recent example of this perceived risk was the widespread fear of a tariff war during the 2008 financial crisis, which lead to the implementation of monitoring systems described in Bown (2011) and pledges by G-20 countries not to 'repeat the historic mistakes of protectionism of previous eras'. Ex-post the worst fears have not be realized, perhaps because there was a network of TAs in place, but what matters more is the perceived ex-ante risk of protectionist spikes.[6]

Focusing on the second type of trade-policy uncertainty, it is easier to examine empirically its significance. Since the relevance of our theory rests more on the presence of uncertainty in *noncooperative* trade policies, it seems natural to focus on volatility in trade policy before the inception of GATT. One of the few available measures that are comparable over time and countries in this period is the import-weighted tariff. Figure 1 presents the distribution of this measure for several countries in 1865-1913. The figure visually suggests a considerable amount of volatility in most of these countries.[7] To place this trade policy volatility into perspective, we can compare it to

[5]Note that our model can accommodate both types of policy uncertainty: the high-frequency/small-magnitude type would be captured by significant probability weight on policy levels relatively close to the mean, while the infrequent/large-magnitude type would be captured by a small probability weight on policy levels that are much higher than the mean level (and of course the two types of risk can co-exist in the same distribution).

[6]This point is made in Handley and Limão (2012). More on this paper below.

[7]In particular, we note that a number of countries have large inter-quartile ranges (represented by the edges of the boxes), up to about 14 percentage points (Russia), and several countries have unusually high periods of protection relative to their typical levels (e.g. the UK, US, Argentina and Brazil had tariffs exceeding their 3rd quartile by more than 1.5 times the inter-quartile range). A legitimate question is whether the measured volatility of the average tariff for a country over time

that of the (log) terms-of-trade. The typical standard deviation for the (log) tariff across countries is .026, which is about 1/4 the magnitude of the typical standard deviation for the (log) terms-of-trade. This is a reasonably large value given that the terms-of-trade of 2/3 of the countries in this sample reflect commodity export prices, which are known to be very volatile.[8]

[**FIGURE 1 HERE**]

There is also some evidence that trade-policy volatility decreases after countries sign TAs. For example, Figure 2 plots the US average tariffs until 1961. The standard deviation of the policy before 1934 is at least twice as high as during 1934-61 – a period marked by the Reciprocal Trade Agreement Act (1934) and the signing of GATT (1948).[9] At the anecdotal level, the higher volatility of US trade policy before 1934 is a reflection not only of the well-known Smoot-Hawley tariff hike of 1930, but also of several major changes in the tariff code (see Irwin, 1998, for a detailed account of these episodes).[10] To be clear, the fact that non-cooperative tariffs are more volatile than cooperative tariffs does not necessarily imply that there is an uncertainty-reducing *motive* for TAs in the sense defined in this paper. This is because our notion of uncertainty-reducing motive is based on a counterfactual agreement (the mean-preserving agreement) that is fully contingent and preserves

---

reflects volatility at the tariff-line level or might instead be driven by volatility in the import weights used to construct average tariffs. It seems unlikely to us that the volatility in average tariffs is driven mainly by changes in import weights for a country over time, for a couple of reasons. First, to replicate the observed volatility in average tariffs, one would need large shifts in import weights within short time periods, and this would be counterfactual because a country's import basket is typically diversified and adjusts slowly over time. Second, at the anecdotal level, in the period we consider, some of the changes in tariffs are clearly due to statutory changes in policy. For example, the US average tariff in the pre-GATT era reflected changes in statutory tariffs, some of which were due to switches in the party in control of Congress (as discussed by Irwin, 1998).

[8] Our calculation of the terms-of-trade volatility is based on data from Blattman, Hwang and Williamson (2007). They measure short-term volatility using the standard deviation of the detrended terms-of-trade. If we do the same (i.e. apply a Hodrick-Prescott filter), we find that the relative short run volatility of the tariff to terms-of-trade is 1/4, which is similar to the one using the underlying series.

[9] Even if some of the trends of protection and liberalization were anticipated and thus not uncertain we are still left with short-term volatility. Using a standard Hodrick-Prescott filter, the short-term volatility prior to 1934 is 2.4 times higher than in the 1934-61 period.

[10] We are not aware of studies that investigate more systematically whether trade agreements reduce policy volatility. Cadot, Olarreaga and Tschopp. (2010) focus on the volatility of agricultural policies, presenting evidence that it was significantly reduced by regional trade agreements. Rose (2004) and Mansfield and Reinhardt (2008) examine the effect of trade agreements on the volatility of trade flows.

the non-cooperative tariff mean, while real-world TAs do not preserve tariff means, and moreover the low volatility of cooperative tariff may simply reflect contracting frictions that cause rigidity in the agreement.

[FIGURE 2 HERE]

Next we discuss the related literature. One strand of literature that is related to our paper focuses on how uncertainty, in conjunction with contracting imperfections, affects the optimal design of TAs. For example, Horn, Maggi and Staiger (2010), Amador and Bagwell (2013) and Beshkar and Bond (2012) show that the presence of uncertainty and contracting imperfections can explain the use of rigid tariff bindings.[11] In contrast to these papers, we focus on the uncertainty-managing motive for a TA and the gains that a TA can provide by regulating policy uncertainty.

Also related is the work by Handley and Limão (2012) and Handley (2014). These papers examine theoretically and empirically the impact that TAs have on trade flows specifically through the channel of removing the risk of future increases in protection. Handley and Limão (2012) find evidence that Portugal's accession to the EC eliminated the risk of exporters losing pre-existing preferences and facing MFN tariffs in the EC. Handley (forthcoming) examines the impact that the imposition of WTO tariff bindings had on exports to Australia, finding evidence that a significant portion of the impact of such bindings on export growth is due to the reduction in the risk of tariff hikes. The contribution of our paper is very distinct. These two papers examine how trade flows are impacted by exogenous reductions in the risk of protection increases, whereas our paper focuses on the endogenous formation of trade policy in non-cooperative and cooperative scenarios, and in particular it examines under what conditions there is "too much" uncertainty in non-cooperative policies, and what are the gains from "correcting" the degree of policy uncertainty through a TA.

The paper is organized as follows. Section II presents our basic reduced-form framework. Section III examines the standard trade model with political economy shocks. Section IV extends the model by allowing for more general economic shocks. Section V presents our "sufficient statistic" approach. Section VI extends the analysis to allow for ex-ante investments. In section VII we consider a setting with two

[11] These contracting imperfections take the form of contracting costs in Horn, Maggi and Staiger; of private information in Amador and Bagwell; and of costly state verification in Beshkar and Bond.

symmetric policy-active countries. Section VIII concludes. The Appendix contains the proofs of our results.

II. Basic framework

To make our points transparent, we start by focusing on a two-country setting where only one country is policy-active, hence there is a one-way international policy externality. In this section we model government objectives in reduced form, as functions of a trade policy and an underlying shock; in the next section we will “open up” the black box of government objectives in the context of a standard trade model.

There are two countries, Home and Foreign. The Home government chooses a trade barrier t, while the Foreign government is passive. We let $G(t, \lambda)$ denote the Home government’s objective function, where λ is interpreted as an exogenous shock to this government’s policy preferences; this could represent for example a politically-adjusted welfare function, with λ a political-economy parameter (e.g. the extra weight attached to a special-interest group) or an economic parameter. We let $F(\lambda)$ denote the c.d.f. of λ. We assume that G is concave in t and satisfies the single crossing property $G_{t\lambda} > 0$. The Foreign government’s objective is $G^*(t, \lambda)$. We assume that an increase in the trade barrier hurts Foreign: $G^*_t < 0$. The governments’ joint payoff is denoted by $G^W(t, \lambda) = G(t, \lambda) + G^*(t, \lambda)$. We assume G^W is concave in t and satisfies the single crossing property $G^W_{t\lambda} > 0$. The role of the single-crossing properties will be apparent shortly.

As we will discuss in the next section, this reduced-form framework can be interpreted as capturing a two-sector, perfectly-competitive world in which a large country trades with a small welfare-maximizing country, and in which a TA is motivated by a terms-of-trade externality.[12] But we note that this framework could also be applied to settings where TAs are motivated by externalities unrelated to terms-of-trade as emphasized by "new trade" models of agreements (e.g. Ossa, 2011, Bagwell and Staiger, 2012, Mrazova, 2011).

We start by describing the non-cooperative policy choice. We assume the Home government observes λ before choosing its trade policy, hence the noncooperative

[12] In the literature on trade agreements there is a small tradition of models with a small country and a large country, a prominent example being McLaren (1997).

policy is given by:

$$t^N(\lambda) = \arg\max_t G(t, \lambda).$$

The single crossing property $G_{t\lambda} > 0$ implies that $t^N(\lambda)$ is increasing. The distribution of the shock $F(\lambda)$ and the shape of the $t^N(\cdot)$ schedule induce a distribution for the noncooperative policy t^N.

We now describe our assumptions regarding the TA. The agreement is signed ex ante, before λ is realized, so the timing is the following: (0) the TA is signed; (1) λ is realized and observed by both countries; (2) t is implemented and payoffs are realized. We assume that λ is verifiable and there are no costs of contracting, so the agreement can be contingent on λ. As we mentioned in the Introduction, given that our main focus is on the *potential* gains from regulating policy uncertainty relative to the noncooperative equilibrium, abstracting from contracting imperfections is a natural first step.[13]

We assume that the TA maximizes the governments' expected joint payoff EG^W,[14] so the (unconstrained) optimal TA is given by

$$t^A(\lambda) = \arg\max EG^W(t, \lambda).$$

The single crossing property $G^W_{t\lambda} > 0$ implies that $t^A(\lambda)$ is increasing.

What motivates governments to sign a TA in this setting is the presence of an international policy externality, which causes the noncooperative policy choice to be inefficient. When we introduce an explicit trade structure in the next section, this externality will operate via terms-of-trade, but for now this can be interpreted as a more general international policy externality.

[13] While our assumption of frictionless contracting serves to focus more sharply on the questions we are addressing, we note that the GATT-WTO does include a number of contingent clauses, for example the "escape clauses" in GATT Articles XIX and XXVIII. For a model that endogenizes the degree to which a trade agreement is contingent, based on the presence of contracting costs, see Horn, Maggi and Staiger (2010).

[14] This implicitly assumes that international transfers are available and that transfers enter governments' payoffs linearly, so that Home's payoff is given by $G + T$ and Foreign's payoff by $G^* - T$, where T is a transfer from Foreign to Home. The transfer can be interpreted for example as a non-trade policy concession that serves as a form of compensation between governments. Focusing on a transferrable-utility setting seems like a natural choice given that we are abstracting from any form of international transaction costs. We also note that the need for government-to-government transfers would be reduced or even eliminated in a more symmetric setting where both countries are policy-active. In section VII we consider a fully symmetric setting, and in such a setting governments select the optimal symmetric agreement, which maximizes the sum of their expected payoffs.

The international policy externality is transmitted through the whole *distribution* of t. For example, if Home's policy schedule $t(\lambda)$ is changed in such a way that the mean of t remains unchanged but the degree of uncertainty in t changes, this will have an impact on Foreign's expected welfare EG^*. In order to isolate the "uncertainty motive" for a TA from the "mean motive", we consider the following thought experiment: if we constrain the TA to keep the average t at the noncooperative level, is there any role left for a TA? This is the idea behind our notion of "mean preserving agreement" (MPA). If the optimal MPA changes the riskiness of t relative to the noncooperative policy $t^N(\lambda)$, we say that there is an uncertainty-*managing* motive for a TA. And in this case, if the optimal MPA decreases (increases) the riskiness of t relative to $t^N(\lambda)$, we say that there is an uncertainty-*reducing* (-increasing) motive for a TA.

Formally, the optimal MPA is defined as

$$t^{MPA}(\lambda) = \arg\max_{t(\lambda)} EG^W(t(\lambda), \lambda) \text{ s.t. } Et(\lambda) = Et^N(\lambda). \tag{II.1}$$

where the operator E denotes an expectation over λ.

Before we study the optimal MPA, we can build intuition by considering a local argument for the simplest possible case. Consider the case where λ affects Foreign only through the policy t, so that its payoff is simply $G^*(t)$. This can be interpreted as a scenario in which λ represents a domestic political-economy shock in the Home country.

Let us start from the noncooperative policy $t^N(\lambda)$ and ask: how can we change the policy schedule locally to achieve an increase in $EG^W = EG + EG^*$, while preserving the mean of the policy? Since $t^N(\lambda)$ maximizes EG, a small change from $t^N(\lambda)$ will have a second-order effect on EG and a first-order effect on EG^*. Clearly, then, to achieve an increase in EG^W we must increase EG^*. Suppose G^* is convex in t: then if we change the policy schedule (slightly) in such a way that the new policy is a mean-preserving spread of $t^N(\lambda)$, this will increase EG^* (by the well-known Rotschild-Stiglitz, 1970, equivalence result) and thus EG^W will also increase. Likewise, if G^* is concave in t, we can achieve an increase in EG^W by making a (slight) mean-preserving compression of $t^N(\lambda)$. Therefore this argument suggests that the key condition determining whether the optimal MPA increases or decreases policy uncertainty is the concavity/convexity of Foreign's objective with respect to t.

Of course, the argument above suggests only a sufficient condition for local improvement over the noncooperative outcome; in particular, one can improve over the noncooperative outcome in many other ways, including by changing the policy schedule in ways that are neither a mean-preserving compression nor spread of $t^N(\lambda)$. But as we show below, this intuition does carry over to the globally optimal MPA in the case of political-economy shocks (when the single-crossing properties are satisfied).

Importantly, however, the Rotschild-Stiglitz type argument no longer applies if the shock λ affects the Foreign payoff G^* directly as well as through the policy t. In this case, it is not enough to know whether Foreign's objective is concave or convex in t to determine how the optimal MPA will change policy uncertainty, as we now show.

To derive the FOCs for the optimal MPA problem in (II.1) we set up the Lagrangian:

$$L = EG^W(t,\lambda) + \psi\left(Et^N(\lambda) - Et(\lambda)\right) \tag{II.2}$$

Since the multiplier ψ is constant with respect to λ, we can rewrite the Lagrangian as follows

$$L = \int [G^W(t,\lambda) + \psi\left(t^N(\lambda) - t(\lambda)\right)]dF(\lambda) \tag{II.3}$$

and since we can maximize this pointwise we obtain the following FOCs

$$\begin{aligned} G_t^W(t(\lambda),\lambda) &= \psi \text{ for all } \lambda \\ Et^N(\lambda) &= Et(\lambda) \end{aligned}$$

Note that the FOC requires the marginal contribution of t to joint surplus, G_t^W, to be equalized across states (realizations of λ), and in particular G_t^W should be equal to the multiplier ψ, which is easily shown to be negative. Also note that the FOC for the unconstrained optimal agreement is given by $G_t^W(t,\lambda) = 0$, so both for the unconstrained optimum and for the optimal MPA, G_t^W is equalized across states, but in the former case it is equalized at zero, while in the MPA case it is equalized at some negative constant.

Using the FOC we can prove:

Lemma 1. *(i) If* $\frac{d}{d\lambda}G_t^*(t^N(\lambda),\lambda) < 0$ *(*> 0*) for all* λ*, then* $t^{MPA}(\lambda)$ *intersects* $t^N(\lambda)$ *once and from above (below). (ii) If* $\frac{d}{d\lambda}G_t^*(t^N(\lambda),\lambda) = 0$ *for all* λ*, then* $t^{MPA}(\lambda) = t^N(\lambda)$ *for all* λ*.*

Figure 3 illustrates Lemma 1 graphically for the case $\frac{d}{d\lambda}G_t^*(t^N(\lambda),\lambda) < 0$. The basic intuition for the result can be conveyed by focusing on the case in which λ can take only two values, say λ^H and λ^L. Let us start from $t^N(\lambda)$ and ask: how can we improve the ex-ante joint payoff? Given the mean-preservation constraint, there are only two ways to modify the schedule $t^N(\lambda)$: decreasing t for $\lambda = \lambda^H$ and increasing t for $\lambda = \lambda^L$ (that is, flattening the schedule), or vice-versa (that is, steepening the schedule). Intuitively it is preferable to reduce t in the state where it is more important to do so, that is where the international externality is stronger (more negative). If $\frac{d}{d\lambda}G_t^*(t^N(\lambda),\lambda) < 0$, then the international externality is stronger in the high-λ state, so it is preferable to flatten the policy schedule relative to $t^N(\lambda)$. Similarly, if $\frac{d}{d\lambda}G_t^*(t^N(\lambda),\lambda) > 0$ the objective can be improved by making the opposite change, that is, steepening the schedule relative to $t^N(\lambda)$. The proof of Lemma 1 (in Appendix) extends this basic logic to the case of continuous λ. Notice that Lemma 1 does not rely on the single crossing properties we assumed for G and G^W, while the next result does.

[FIGURE 3 HERE]

Lemma 1 leads directly to our first proposition. In the proposition, we say that there is an uncertainty-reducing (-increasing) motive for a TA if $t^{MPA}(\lambda)$ is a mean preserving compression (spread) of $t^N(\lambda)$.

Proposition 1. *(i) If* $\frac{d}{d\lambda}G_t^*(t^N(\lambda),\lambda) < 0$ *(*> 0*) for all* λ*, there is an uncertainty-reducing (-increasing) motive for a TA. (ii) If* $\frac{d}{d\lambda}G_t^*(t^N(\lambda),\lambda) = 0$ *for all* λ *then* $t^{MPA}(\lambda) = t^N(\lambda)$*, hence there is no uncertainty-managing motive for a TA.*

Proposition 1 states that the *direction* of the uncertainty-managing motive for a TA, if any, is determined by how the shock λ affects the marginal international externality G_t^* at the noncooperative equilibrium, taking into account its direct effect and its indirect effect through the policy. In particular, if $G_t^*(t^N(\lambda),\lambda)$ is decreasing (increasing) in λ then there is an uncertainty-reducing (-increasing) motive for a TA. Writing $G_t^*(t^N(\lambda),\lambda) = G_{tt}^{*N} \cdot \frac{dt^N}{d\lambda} + G_{t\lambda}^{*N}$ (where we use a superscript N to indicate that a function is evaluated at $t^N(\lambda)$), the uncertainty motive for a TA can be traced to two key determinants: (a) Foreign's *policy-risk preference* (captured by G^*_{tt} and weighted by $\frac{dt^N}{d\lambda}$), and (b) the direct impact of the shock λ on the marginal international externality holding t constant (as captured by $G^*_{t\lambda}$), which we refer to as the *externality-shifting* effect.

Proposition 1 makes clear that the source of uncertainty matters. In particular, we can distinguish between two types of shock: (1) a "political economy" shock, which affects the Foreign country only through the policy t (in which case $G^* = G^*(t)$); and (2) an "economic" shock, which affects the Foreign country not only indirectly through the policy t but also directly (in which case $G^* = G^*(t, \lambda)$).

In the case of "political economy" shocks, Proposition 1 says that the uncertainty motive for a TA is determined solely by Foreign's preference for policy risk, as captured by the sign of G^*_{tt}. This confirms our initial intuition based on Rotschild and Stiglitz's (1970) result: when Foreign's objective is concave in t, a MPS in t reduces EG^*, so there is a negative "policy-risk externality," hence the noncooperative policy is "too risky" (with the reverse logic holding if Foreign's objective is convex in t).[15] In the case of "economic" shocks, on the other hand, Proposition 1 states that Foreign's policy-risk preference (the sign of G^{*N}_{tt}) is no longer sufficient to determine whether there is "too much" or "too little" risk in the noncooperative policy, because the externality-shifting effect ($G^*_{t\lambda}$) comes into play. In this case, the direction of the uncertainty motive for a TA is determined by whether the international policy externality G^*_t is increasing or decreasing in λ at the noncooperative equilibrium.

Before concluding this section, we mention an alternative thought experiment that one could consider to isolate the uncertainty motive for a TA. Suppose the Home government can choose a contingent policy $t(\lambda)$ subject to the constraint that this policy have the same mean as the optimal agreement policy $t^A(\lambda)$. If such "mean-preserving unilateral" policy is more risky than $t^A(\lambda)$, then we say that the noncooperative policy is "too risky", and so there is an uncertainty-reducing motive for a TA. One can show that, under this alternative thought experiment, the direction of the un-

[15] We highlight however that, even in this case where the result is intuitive, it is far from self-evident: a priori the optimal MPA could have entailed any mean-preserving change in t relative to $t^N(\lambda)$, and since the MPS risk criterion is a partial ordering, it was not a priori obvious that the distribution of $t^{MPA}(\lambda)$ could be ranked in a MPS sense relative to that of $t^N(\lambda)$. We also note here that if $G^*_{tt} < 0$ (> 0) then the optimal MPA policy is a "simple" mean preserving spread (compression) of the noncooperative policy, meaning that the respective cdf's cross only once (as shown in the proof of Proposition 1).

certainty motive is again determined by the sign of $\frac{d}{d\lambda}G_t^*$, but this time evaluated at $t^A(\lambda)$ rather than at $t^N(\lambda)$.[16] As a consequence, if λ is a "political economy" shock, the two thought experiments yield the same answer (there is an uncertainty-reducing motive for a TA if and only if $G_{tt}^* < 0$). If λ is an "economic" shock, on the other hand, both thought experiments indicate that the uncertainty motive depends on Foreign's policy-risk preference (G_{tt}^*) and on the externality-shifting effect ($G_{t\lambda}^*$), but the relative weight of these two terms differs (in one case G_{tt}^* is weighted by $t^{N'}(\lambda)$, in the other case it is weighted by $t^{A'}(\lambda)$). In what follows we base our analysis on the MPA thought experiment. The main reason is that, as we will show later, focusing on the MPA allows us to characterize the gains from regulating trade-policy uncertainty and trade-policy mean in terms of quantities that can in principle be observed or estimated, while the alternative thought experiment does not share this property.[17]

A. Gains from regulating policy uncertainty and policy mean

In this section we develop a simple formula for the gains from regulating policy uncertainty relative to the more standard gains from regulating the policy mean. We will later apply this formula in the context of our trade model, in order to analyze how the relative gains from regulating policy uncertainty depend on economic fundamentals, and to illustrate how they can potentially be quantified with data.

Ideally one would focus on the gain from the optimal MPA, that is the increase in EG^W associated with a move from $t^N(\lambda)$ to $t^{MPA}(\lambda)$. Here we focus on a simpler task, namely, evaluating the gain from a *small* improvement in policy uncertainty. In particular, we consider a small mean-preserving change in the policy schedule starting from $t^N(\lambda)$ and evaluate the effect of this change on EG^W. We will then evaluate the gain from a small reduction in the policy mean, and finally derive an expression for the ratio between the two gains. Our approach of focusing on small policy changes is similar in spirit to the analysis of "piecemeal" policy reforms commonly applied in second-best theory.

Consider moving from $t^N(\lambda)$ to $t^N(\lambda) - \delta(t^N(\lambda) - \bar{t}^N)$, where δ is a small constant and $\bar{t}^N$ is the mean of $t^N(\lambda)$. Clearly, if $\delta > 0$ ($\delta < 0$) this represents a small

[16] The proof of this statement is available upon request.

[17] To be more specific, with our MPA thought experiment we approximate the gains from regulating trade-policy uncertainty starting from the noncooperative tariff, which is in principle observable. Under the alternative thought experiment the starting point for the approximation would be the "mean-preserving unilateral" tariff choice (defined above), which is unobservable.

mean-preserving compression (spread) of $t^N(\lambda)$. The resulting change in EG^W can be approximated as follows:

(II.4)
$$\frac{\partial EG^W(t^N(\lambda)-\delta(t^N(\lambda)-\bar{t}^N),\lambda)}{\partial\delta}\bigg|_{\delta=0} = -E[G_t^*(t^N,\lambda(t^N))(t^N-\bar{t}^N)]$$

$$\approx -E\left[\left(G_t^*(\bar{t}^N,\lambda(\bar{t}^N))+\frac{dG_t^*(t^N,\lambda(t^N))}{dt^N}\bigg|_{t^N=\bar{t}^N}\cdot(t^N-\bar{t}^N)\right)(t^N-\bar{t}^N)\right]$$
$$= -\frac{dG_t^*(t^N,\lambda(t^N))}{dt^N}\bigg|_{t^N=\bar{t}^N}\cdot\sigma^2_{t^N}$$

In the first line of (II.4) we use the fact that $G_t = 0$ at the noncooperative policy, and employ a change of variables from λ to t^N, letting $\lambda(t^N)$ denote the inverse of $t^N(\lambda)$ (with the expectation now taken with respect to t^N). In the second line we use a first-order Taylor approximation of $G_t^*(t^N,\lambda(t^N))$ around $\bar{t}^N$.

The last line of (II.4) states that the effect of a small change in policy uncertainty on EG^W is the product of two components. The first one is analogous to the derivative $\frac{dG_t^*(t^N(\lambda),\lambda)}{d\lambda}$, except for the change of variable from λ to t^N. Recall from Proposition 1 that the sign of this derivative determines the direction of the uncertainty-managing motive: if the international externality G_t^* is stronger when the noncooperative policy is higher, there is value to reducing policy uncertainty. The second component is the variance of t^N, which intuitively magnifies the value of managing policy uncertainty.

Since the sign of δ can be chosen to ensure a positive gain, we can write the approximate value of a small change in policy uncertainty as

(II.5)
$$\tilde{V}^{MPA} \equiv \left|\frac{dG_t^*(t^N,\lambda(t^N))}{dt^N}\right|_{t^N=\bar{t}^N}\cdot\sigma^2_{t^N}$$

Next we focus on the gains from regulating the mean level of the policy. A natural approach is to define an "uncertainty-preserving agreement" (UPA) in the following way. Consider a parallel downward shift of the $t^N(\lambda)$ schedule, $t^N(\lambda)-\kappa\bar{t}^N$, where κ is a positive constant. This shift reduces the mean of the policy by a factor κ but preserves all its central higher moments (variance, skewness, kurtosis), so it is natural to interpret such a shift as one that changes the policy mean while preserving policy uncertainty.

Following similar steps as above, we can approximate the value of a "small" UPA as

$$\text{(II.6)} \qquad \frac{\partial EG^W(t^N(\lambda)-\kappa\bar{t}^N,\lambda)}{\partial\kappa}|_{\kappa=0} = -EG_t^*(t^N,\lambda(t^N))\cdot\bar{t}^N$$

$$\begin{aligned}
&\approx \; -E\left(G_t^*(\bar{t}^N,\lambda(\bar{t}^N)) + \frac{dG_t^*(t^N,\lambda(t^N))}{dt^N}\bigg|_{t^N=\bar{t}^N}\cdot(t^N-\bar{t}^N)\right)\cdot\bar{t}^N \\
&= \; -G_t^*(\bar{t}^N,\lambda(\bar{t}^N))\cdot\bar{t}^N \equiv \tilde{V}^{UPA}
\end{aligned}$$

Intuitively, the gain from reducing the mean policy level is approximately equal to the marginal international externality from the policy $(G_t^*(\cdot))$ evaluated at certainty and scaled up by the mean policy level.

Next we write down the gains from a small change in policy uncertainty versus a small change in policy mean:

$$\text{(II.7)} \qquad \frac{\tilde{V}^{MPA}}{\tilde{V}^{UPA}} = \left|\frac{d\ln G_t^*(t^N,\lambda(t^N))}{dt^N}\right|_{t^N=\bar{t}^N}\cdot\frac{\sigma^2_{t^N}}{\bar{t}^N}$$

Notice that $\frac{\tilde{V}^{MPA}}{\tilde{V}^{UPA}}$ can be interpreted as the value of a 1 percent change in the standard deviation of t relative to a 1 percent reduction in the level of t, starting from the non-cooperative equilibrium.[18] Later on, when we apply this framework to a simple competitive trade model, we will examine how the ratio $\frac{\tilde{V}^{MPA}}{\tilde{V}^{UPA}}$ depends on underlying model parameters. Before proceeding, though, it is important to be clear on the limitations of this approach: this ratio captures the gains from a "small" MPA relative to a "small" UPA, rather than the gains from the optimal MPA relative to the optimal UPA. At the same time, intuitively the former ratio should be informative about the latter. In particular, when we examine how $\frac{\tilde{V}^{MPA}}{\tilde{V}^{UPA}}$ depends on model parameters, a reasonable conjecture is that the qualitative results would not be overturned if one were to consider the "ideal" measure of relative gains. In what follows, with a slight abuse of terminology, we will refer to $\frac{\tilde{V}^{MPA}}{\tilde{V}^{UPA}}$ as the "relative gains from regulating policy uncertainty".

[18] To understand this, note that the standard deviation of $t^N - \delta(t^N - \bar{t}^N)$ is equal to $(1-\delta)$ times the standard deviation of t^N.

Finally, one can consider the value of a small joint improvement in policy mean and policy uncertainty starting from the noncooperative equilibrium. Clearly, this value is given by a weighted average of expressions (II.4) and (II.6) above, with the weights determined by the relative change in policy mean and uncertainty. Below we will apply this observation in the context of our economic structure.

III. The uncertainty motive for a TA in a competitive trade model

A. Setup

We now impose more structure on the model in order to examine how the uncertainty motive for a TA depends on economic fundamentals. We consider a standard two-country, two-good trade model with competitive markets. We assume Home is the natural exporter of the numeraire good, indexed by 0, while Foreign (the small country) is the natural exporter of the other good, which has no index.

Let p (resp. p^*) denote the price of the nonnumeraire good in Home (resp. Foreign). We will often use the logarithms of prices, letting $\pi \equiv \ln p$ and $\pi^* \equiv \ln p^*$. The Home country can choose an ad-valorem tariff on imports of the non-numeraire good. Let $t \equiv \ln \tau$, where τ is the ad-valorem tariff factor. We also allow for an exogenous iceberg trade cost and denote the logarithm of this cost factor by γ. The reason we allow for trade costs is not only that such costs are important empirically, but because they will play an important role in determining the gains from regulating policy uncertainty, as will become clear below. By the usual arbitrage condition, if the tariff is not prohibitive then we must have $\pi^* = \pi - t - \gamma$. Since Foreign has no policy of its own, we can refer to π^* as Foreign's "terms-of-trade" (TOT). Since Foreign is small, π is determined entirely in the Home country, so we can leave the market clearing condition that determines π in the background.

The reason we use the logarithms of relative prices is the following. In general equilibrium settings with uncertainty about relative prices, the standard notion of arithmetic mean preserving spread leads to results that are sensitive to the choice of numeraire, as pointed out by a number of papers, for example Flemming, Turnovsky and Kemp (1977). These papers have argued that a more robust approach is to define an increase in relative-price risk as a *geometric* mean preserving spread (GMPS) of the relative price, that is an arithmetic mean preserving spread of the log of the relative price. For this reason, we will follow the mainstream literature and consider arithmetic mean preserving changes in $\ln \tau$ (and hence in $\ln p^*$), so our results are not sensitive to the choice of numeraire.[19]

We next impose some standard assumptions on preferences and technology. To make the key points we only need to specify the economic structure in the Foreign country. On the technology side, we assume constant returns to scale with a strictly concave PPF, so that supply functions are strictly increasing. This allows us to describe the supply side through a GDP (or revenue) function. Letting p^* be the domestic relative price and (q_0^*, q^*) the outputs, we define $R^*(p^*) \equiv \max_{q_0^*, q^*}\{q_0^* + p^* q^*\}$ s.t. $(q_0^*, q^*) \in Q^*$, where Q^* is the set of feasible outputs.

On the preference side, we assume that all citizens have identical and homothetic preferences. This implies that indirect utility takes the form $U\left(\frac{y^*}{\phi^*(p^*)}\right)$, where y^* is income in terms of numeraire and $\phi^*(p^*)$ a price index. It is natural to refer to $\frac{y^*}{\phi^*(p^*)}$ as the representative individual's "real income". For the purposes of comparative statics it is convenient to parametrize the degree of risk aversion, so we assume that $U(\cdot)$ exhibits constant relative risk aversion (CRRA), indexed by the parameter θ.

All citizens have identical factor endowments, and the population measure is normalized to one. There are no international risk-sharing markets, so that the Foreign country cannot diversify away its income risk.[20] The Foreign government maximizes social welfare, so we can write[21]

[19] To elaborate on this point, the way one defines a change in risk is an axiomatic choice: there is no single, "correct" way to define a change in risk, so it seems reasonable to choose a definition that satisfies certain desirable properties (axioms) in the relevant context of analysis. In a general equilibrium model such as ours, the GMPS notion is a reasonable choice for the reason discussed in the text. The same approach is followed by Eaton (1979), which we discuss below.

[20] International risk sharing is arguably very incomplete in reality, and our assumption captures this incompleteness in an extreme way, but even if had perfect risk-sharing markets in the model, in general there would still be an uncertainty motive for a TA: in this case, the model would be equivalent to one where agents are income-risk neutral, and as we show below, there would typically be an uncertainty-increasing motive for a TA.

[21] We note that the assumption of risk-averse citizens is not in contradiction with the assumption – discussed in footnote 14 – that the government's utility is transferrable. Recall that the Foreign government's payoff is assumed to be $G^* - T$, where G^* is the utility of the representative citizen and T the transfer made to the Home government (e.g. in the form of a non-trade policy concession). We view the assumption of transferrable government utility as a convenient modeling device that allows us to focus on the TA that maximizes the governments' joint payoff. A more restrictive assumption that is implicit in our setting, on the other hand, is that the TA cannot specify contingent transfers that can in turn be used to provide insurance to citizens: if this were the case, a TA could be used as an international risk-sharing mechanism, thus making risk aversion irrelevant. Contingent transfers

$$G^* = \frac{1}{\theta}\left(\frac{R^*(p^*)}{\phi^*(p^*)}\right)^\theta$$

We first focus on the case in which shocks are of the "political economy" type, that is, shocks originate in Home and affect Foreign only through t. Subsequently we extend the analysis to the case of more general "economic" shocks.

Given this simple structure, we do not have to be explicit about Home's technology and preferences, so we will keep them in the background. All that matters for our purposes is Home's non-cooperative tariff schedule $t^N(\lambda)$. Finally, we assume that the trade pattern cannot switch as a result of the shock, that is, Foreign exports the nonnumeraire good for all values of λ in its support.

The analysis of section II shows that the key to gauge the uncertainty motive for a TA is to consider how the marginal international externality exerted by the Home tariff, G_t^*, responds to the shock λ. In our model, Home's tariff exerts only a TOT externality on Foreign welfare, which is given by

$$G_t^* = -v^{*\theta} \cdot \Omega^* \tag{III.1}$$

where $v^* = \frac{R^*}{\phi^*}$ is Foreign's real income and $\Omega^* \equiv \frac{p^* x^*}{R^*}$ is Foreign's degree of openness (export share of GDP). Intuitively, the degree of openness Ω^* captures the impact of an increase in t on Foreign's real income through TOT, and the factor $v^{*\theta}$ is related to the marginal utility of income: with $\theta < 0$, the externality is stronger when real income (v^*) is lower (for a given level of openness), because the marginal utility of income is higher. In what follows we will refer to $v^{*\theta}\Omega^*$ as the "adjusted" degree of openness.[22]

between governments can be allowed in our model only if they take a non-monetary form so that they cannot be used to provide insurance to citizens.

[22] As stated earlier, in this model the underlying motive for a TA is the presence of a TOT externality. To be more precise, the reason why the noncooperative equilibrium is inefficient is not the presence of a TOT externality *per se*, but the fact that the Home country has monopoly power over TOT. To confirm this point, consider an alternative version of this model where the Home country is replaced by a continuum of symmetric small countries (all affected by a common λ shock): in such a setting it can be verified that the noncooperative equilibrium would be efficient for all λ.

B. Income-risk neutrality.

Since we are adopting the GMPS notion of risk, it is natural to define risk neutrality as indifference with respect to a GMPS of real income, which corresponds to the case: $U(\cdot) = \ln(\cdot)$, or $\theta \to 0$ in the CRRA specification. Thus the government's objective is $G^* = \ln\left(\frac{R^*(p^*)}{\phi^*(p^*)}\right)$, and the international externality is simply $G^*_t = -\Omega^*$.

The key step to apply Proposition 1, given that λ is a political economy shock, is to examine the Foreign country's attitude toward policy risk, as captured by G^*_{tt}. This is given by the impact of t on openness, which is easily shown to be

$$G^*_{tt}|_{\theta\to 0} = \Omega^* \left(\varepsilon^*_x + D^*\right),$$

where ε^*_x is the export supply elasticity and $D^* \equiv 1 - \frac{p^* q^*}{R^*}$ is the import-competing sector share of GDP, which can be interpreted as the degree of income diversification.[23]

We will assume throughout that ε^*_x is nonnegative.[24] Given this assumption, it follows that $G^*_{tt}|_{\theta\to 0} > 0$: thus, in the case of income-risk neutrality, the Foreign country benefits from an increase in policy risk. The intuition for this result is that, since production and consumption can be optimized after observing prices, both the producers' revenue function and the consumers' indirect utility functions (given income) are convex in prices.[25] The insight that a small country may gain from TOT risk in itself is not new to our model, and was pointed out for example by Eaton (1979);[26] what is new is that in light of Proposition 1, the convexity of G^* with respect to t implies that the optimal MPA increases trade-policy uncertainty.

To summarize the discussion thus far, *if individuals are income-risk neutral, there is an uncertainty-managing motive for a TA, but this calls for an increase – rather than a decrease – in trade-policy uncertainty.*

[23] Note that $D^* = 0$ when the country is completely specialized, and $D^* = 1/2$ when the two sectors have equal GDP shares. In interpreting D^* as an index of diversification we are implicitly assuming that the GDP share of the export sector is at least 1/2, so that D^* cannot exceed 1/2.

[24] There is considerable empirical evidence that this is the case in reality for most sectors and most countries (see for example Tokarick, 2010).

[25] It is important to note that this feature extends well beyond the simple perfectly-competitive setting we are considering here. In particular, one might wonder whether the presence of imperfect competition or irreversible investments might make exporting firms' profit functions concave in prices, but even in these circumstances profit functions can be convex in prices. The intuitive reason is that profit functions are convex whenever firms can make *any* ex-post adjustment in their production decisions after observing prices, and this feature is extremely general.

[26] See also Anderson and Riley (1976), who examine how the degree of specialization of a small economy affects its gains from TOT fluctuations.

Evidently, then, if one wants to make economic sense of the WTO-type informal arguments discussed in the introduction, which state that one of the goals of TAs is to reduce trade policy uncertainty, one must depart from the benchmark case of income-risk neutrality in this standard model and focus on the case of income-risk aversion, which is what we do next.

C. Income-risk aversion

Let us now re-examine the Foreign country's preference for trade-policy risk allowing for income-risk aversion ($\theta < 0$).[27] Recalling that the international externality from the tariff is given by $G_t^* = -v^{*\theta} \cdot \Omega^*$ and differentiating this expression with respect to t, we obtain

$$G_{tt}^* = v^{*\theta}\Omega^* \left(\theta\Omega^* + \varepsilon_x^* + D^*\right). \tag{III.2}$$

This expression (derived in Appendix within the proof of Proposition 2), together with the result of Proposition 1, leads to:

Proposition 2. *There is an uncertainty-reducing (-increasing) motive for a TA if $\theta\Omega^* + \varepsilon_x^* + D^* < 0$ (> 0) at the noncooperative equilibrium.*

There are several aspects of Proposition 2 that are worth highlighting. First, if income-risk aversion is sufficiently strong relative to the other parameters of the model (namely if $\theta < -\frac{\varepsilon_x^* + D^*}{\Omega^*}$), then there is an uncertainty-reducing motive for a TA. While the role of risk aversion is quite intuitive, the impact of the other variables – which we focus on next – is more subtle.

Proposition 2 states that, for a given degree of risk-aversion $\theta < 0$, the uncertainty motive for a TA is more likely to be in the direction of reducing policy uncertainty

[27] Note that, even with income risk aversion, in the Foreign country there is still no motive for trade protection, so our assumption that this country practices free trade continues to be without loss of generality given the representative-citizen assumption. As Eaton and Grossman (1985) made clear, in a small country an insurance motive for trade protection can arise only if citizens have heterogenous incomes, at least ex-post. In our setting, Foreign citizens are always homogenous, even ex-post. Here we also mention the paper by Young and Anderson (1982), who compare the effects of quotas and tariffs for a small economy where individuals are risk averse and the government faces political-economy constraints (in the form of a minimum expected output level). We note that the focus of this paper is very different from ours, since it focuses on optimal policies for a small country, while in our setting the optimal policy for the small country is free trade, and our focus is on the externality that the large country's policy fluctuations exert on the small country, and how a TA can correct such externality.

when: (a) the economy is more open (Ω^* is higher); (b) the export supply elasticity ε_x^* is lower; and (c) the economy is more specialized (D^* is lower).[28]

Focus first on the degree of openness Ω^*. This variable affects the uncertainty motive through its interaction with the income-risk preference parameter θ, so the role of openness is in essence to magnify the impact of the citizens' income-risk preference.

Next consider the role of the export supply elasticity ε_x^*. Intuitively, a country that can easily adjust production and consumption as a result of the shocks (that is, a country with a higher ε_x^*) is more likely to have a welfare function that is convex in the foreign tariff, and hence is less likely to benefit from a decrease in tariff uncertainty. This in turn suggests an interesting implication. At the empirical level, lower-income countries tend to have lower export supply elasticities, and this in turn implies that the uncertainty-reducing motive for a TA should be more important for lower-income countries than for higher-income countries.[29]

Focus next on the degree of diversification, D^*. Proposition 2 indicates that, other things equal, the uncertainty motive for a TA is more likely to be in the direction of reducing policy uncertainty if the Foreign country is less diversified. A related remark is the following: assuming that preferences are Cobb-Douglas and the supply function $q^*(p^*)$ is differentiable, if the economy is sufficiently specialized (D^* is sufficiently close to zero) then there is an uncertainty-reducing motive for any $\theta < 0$.[30] Interestingly,

[28] Here we can make the statements in the text a bit more precise. First, when we say that the uncertainty motive is "more likely" to be in the direction of reducing policy uncertainty when a variable x is higher, we mean that as x increases the sign of G_{tt}^{*N} can switch from negative to positive but not vice-versa. Second, in the text we talk about changes in Ω^*, D^* and ε_x^* as if these variables were exogenous, but of course they are not. To make our statements more precise, let $\boldsymbol{\xi}$ denote the vector of all technology and preference parameters (excluding θ). We can think of the key endogenous variables Ω^*, D^* and ε_x^* as functions of $\boldsymbol{\xi}$. Note that θ does not affect these variables. Next note that $\Omega^* \in [0,1]$ and $D^* \in [0,1]$, while $\varepsilon_x^* \geq 0$ by assumption. In the text, when we refer to a change in an endogenous variable, we mean that the parameter vector $\boldsymbol{\xi}$ is being changed in such a way that the variable of interest changes while the others do not. If we include in $\boldsymbol{\xi}$ the whole technology and preference structure, by varying $\boldsymbol{\xi}$ we can span the whole feasible range of Ω^*, D^* and ε_x^*, so this "all else equal" thought experiment can be performed.

[29] See for example Tokarick (2014), who estimates that the median export supply elasticity is 0.52 for low income countries, 0.77 for low/medium income countries, 0.83 for medium/high income countries, 0.92 for high income non-OECD countries, and 1.14 for high income OECD countries. These estimates are based on a standard trade model for a small economy with one export, one import and one non-traded good, with no own consumption of the export good.

[30] To see this, recall that there is an uncertainty-reducing motive if $\theta < -\frac{\varepsilon_x^* + D^*}{\Omega^*}$. In the limit as the country becomes fully specialized, $\frac{p^* q^*}{R^*} \to 1$, hence $D^* \to 0$. Next note that $\varepsilon_x^* = \frac{q^*}{x^*}\varepsilon_q^* - \frac{c^*}{x^*}\varepsilon_c^*$, where ε_q^* is the elasticity of $q^*(p^*)$ and ε_c^* is the elasticity of $c^*(p^*)$. Cobb-Douglas preferences imply $c^* = \frac{\alpha R^*}{p^*}$, where α is the consumption share of the non-numeraire good, hence $\varepsilon_c^* = \frac{d \ln R^*}{d \ln p^*} - 1$;

these twin observations go in the same direction as the one we made above about ε_x^*: to the extent that lower-income countries are more likely to be specialized, our model predicts that the uncertainty-reducing motive for a TA should tend to be more important for lower-income countries.

One way to summarize the discussion above is that the direction of the uncertainty motive for a TA is determined by an overall tradeoff between risk aversion, which operates through the term $\theta\Omega^*$ and pushes toward an uncertainty-reducing motive, and the degree of flexibility of the economy, which is captured by $(\varepsilon_x^* + D^*)$ and pushes toward an uncertainty-increasing motive.

Finally, it is interesting to consider the impact of the exogenous trade cost γ. We consider the following thought experiment: letting γ^{prohib} denote the level of γ for which which there is no trade ($\Omega^* = 0$), we examine the effect of decreasing γ from γ^{prohib} to zero.

Suppose risk aversion is sufficiently strong that in the absence of trade costs ($\gamma = 0$) there is an uncertainty-reducing motive for a TA, that is $\theta < \left(-\frac{\varepsilon_x^* + D^*}{\Omega^*}\right)_{\gamma=0}$.[31] Clearly, as γ drops below γ^{prohib}, initially the uncertainty motive for a TA goes in the direction of increasing policy uncertainty (because $\theta\Omega^*$ is negligible and hence dominated by $\varepsilon_x^* + D^*$), but as γ drops further, the direction of the uncertainty motive will at some point reverse and call for a reduction in policy uncertainty. Thus we can state:

Remark 1. *Assume risk aversion is sufficiently strong, in the sense that* $\theta < \left(-\frac{\varepsilon_x^* + D^*}{\Omega^*}\right)_{\gamma=0}$. *If the trade cost* γ *is close enough to its prohibitive level, there is an uncertainty-increasing motive for a TA* ($\theta\Omega^* + \varepsilon_x^* + D^* > 0$)*, while if* γ *is close enough to zero there is an uncertainty-reducing motive for a TA* ($\theta\Omega^* + \varepsilon_x^* + D^* < 0$)*.*

---

but $\frac{d \ln R^*}{d \ln p^*} = \frac{p^* q^*}{R^*} \to 1$, hence $\varepsilon_c^* \to 0$. Given the assumption that $q^*(p^*)$ is smooth, in the limit as the economy becomes fully specialized clearly $q^{*'}(p^*)$ must approach zero (because of the resource constraint), hence $\varepsilon_q^* \to 0$, which implies $\varepsilon_x^* \to 0$. And since $\Omega^* > 0$, then $\frac{\varepsilon_x^* + D^*}{\Omega^*} \to 0$. So we can conclude that for any fixed $\theta < 0$ the condition $\theta < -\frac{\varepsilon_x^* + D^*}{\Omega^*}$ is satisfied if the economy is sufficiently specialized.

[31] It may be interesting to note that in our quantification exercise for the US-Cuba case (see section V) this condition seems to be largely satisfied. In that context we find that, at the factual level of trade cost γ, there is an uncertainty-reducing motive for a TA, given the available estimates of θ. The condition should then be satisfied a fortiori at the counterfactual level $\gamma = 0$, since lowering γ increases openness Ω^*.

Remark 1 suggests two broad implications of the model, one concerning the evolution of the uncertainty motive for TAs over time and one of a cross-sectional nature. First, as trade costs fall over time, the model suggests that uncertainty-reducing motives for TAs are increasingly likely to emerge, provided citizens are sufficiently risk-averse. And second, since trade costs tend to increase with geographical distance, the model suggests that uncertainty-reducing motives for TAs are more likely to be present (other things equal) for countries within a region.

D. Gains from regulating trade-policy uncertainty

In this section we apply the formulas developed in section A to examine the gains from regulating trade-policy uncertainty relative to its mean.

Given that the political economy shock λ affects Foreign welfare only through Home's tariff t, we have $\frac{dG_t^*(t^N,\lambda(t^N))}{dt^N} = G_{tt}^*(t^N)$. Plugging the expressions for G_t^* and G_{tt}^* in the formulas of section A, we obtain:

Proposition 3. *(i) The value of a small change in tariff uncertainty is* $\tilde{V}^{MPA} = |\theta\Omega^* + \varepsilon_x^* + D^*| \cdot \left(v^{*\theta}\Omega^*\right) \cdot \sigma_{t^N}^2$*; (ii) the value of a small reduction in the tariff mean is* $\tilde{V}^{UPA} = v^{*\theta}\Omega^* \cdot \bar{t}^N$*; and (iii) the relative value of regulating tariff uncertainty is* $\tilde{V}^{MPA}/\tilde{V}^{UPA} = |\theta\Omega^* + \varepsilon_x^* + D^*| \cdot \frac{\sigma_{t^N}^2}{\bar{t}^N}$*, where all expressions are evaluated at the noncooperative equilibrium.*

It is worth highlighting the role of two key determinants of $\tilde{V}^{MPA}$ and $\tilde{V}^{UPA}$: the variance of the noncooperative tariff, $\sigma_{t^N}^2$, which can be interpreted as capturing the degree of uncertainty in the trade policy environment, and the exogenous trade cost γ.

Focus first on the role of $\sigma_{t^N}^2$. Other things equal, an increase in $\sigma_{t^N}^2$ leads to an increase in $\tilde{V}^{MPA}$, while leaving $\tilde{V}^{UPA}$ unaffected. Thus, the gain from a joint improvement in tariff uncertainty and tariff mean – which is given by a weighted average of $\tilde{V}^{MPA}$ and $\tilde{V}^{UPA}$, as discussed in section A – is increasing in $\sigma_{t^N}^2$. Thus, our model suggests that governments should have stronger incentives to sign trade agreements when the trading environment is more uncertain.[32]

[32] While $\tilde{V}^{MPA}$ and $\tilde{V}^{UPA}$ capture only the gains from small policy changes, we can make a similar point by focusing on the gains from the optimal TA, if we take quadratic approximations of the payoff functions. The value of the optimal TA is given by $E[G^W(t^A(\lambda),\lambda) - G^W(t^N(\lambda),\lambda)]$. Consider a mean preserving spread of λ, which captures an increase in underlying uncertainty.

Next we focus on the impact of the trade cost γ, and in particular on how it affects the relative gains from regulating policy uncertainty ($\tilde{V}^{MPA}/\tilde{V}^{UPA}$). We continue to assume $\theta < \left(-\frac{\varepsilon_x^*+D^*}{\Omega^*}\right)_{\gamma=0}$, as in the previous section. As we observed above, there exists a critical level of γ, say $\hat{\gamma}$, for which $\theta\Omega^* + \varepsilon_x^* + D^* = 0$. To simplify, we assume that $\hat{\gamma}$ is unique. Under this assumption, the ratio $\tilde{V}^{MPA}/\tilde{V}^{UPA}$ is non-monotonic in γ, with a minimum value of zero at $\gamma = \hat{\gamma}$. To see this, note that when γ is close to γ^{prohib}, the relative gain is strictly positive (with the gains from the MPA coming from an increase in uncertainty); when γ is equal to $\hat{\gamma}$ the ratio $\tilde{V}^{MPA}/\tilde{V}^{UPA}$ reaches zero; and if γ is lower than $\hat{\gamma}$ this ratio is strictly positive again, but this time the gains from the MPA come from a decrease in uncertainty. Thus we can state:

Remark 2. *Assume that $\theta < \left(-\frac{\varepsilon_x^*+D^*}{\Omega^*}\right)_{\gamma=0}$ and $\hat{\gamma}$ is unique. Then $\tilde{V}^{MPA}/\tilde{V}^{UPA}$ is non-monotonic in γ, with a minimum value of zero at $\gamma = \hat{\gamma}$.*

This result, which can be seen as complementing the result in Remark 1, suggests that, if trade costs fall over time, the relative gains from regulating trade-policy uncertainty may initially decrease, but should eventually grow in magnitude, provided citizens are sufficiently risk averse.

E. Impact of policy uncertainty on trade volume

The next question we address is, what is the impact of the optimal MPA on the expected volume of trade? To fix ideas, suppose that the optimal MPA leads to a mean preserving compression in t. Writing trade volume as $x^*(\pi^*)$, the change in expected log trade due to the MPA is $\int \ln x^*(\pi^*)\, d\left(F_{MPA}(\pi^*) - F_N(\pi^*)\right)$, where $F_N(\pi^*)$ (resp. $F_{MPA}(\pi^*)$) is the distribution of π^* induced by $t^N(\lambda)$ (resp. $t^{MPA}(\lambda)$). Noting that a mean preserving compression in t leads to a mean preserving compression in π^*, by standard Rotschild-Stiglitz logic it is immediate to conclude that expected log trade increases if and only if Foreign's export supply elasticity ε_x^* is decreasing in π^*. Also note that the same conclusion applies to the (log) trade value $\pi^* + \ln x^*(\pi^*)$, since an MPA keeps $E(t)$ and thus $E(\pi^*)$ unchanged.

This will increase the value of the TA if and only if $G^W(t^A(\lambda),\lambda) - G^W(t^N(\lambda),\lambda)$ is convex in λ. Assuming that all third derivatives of G and G^W are zero, this is the case if $G_{tt}^W\left(t^{A'}\right)^2 + 2G_{t\lambda}^W t^{A'} - \left(G_{tt}^W\left(t^{N'}\right)^2 + 2G_{t\lambda}^W t^{N'}\right) > 0$. Using $t^{A'} = \frac{G_{t\lambda}^W}{-G_{tt}^W}$ and simplifying, this condition becomes $\left(t^{A'} - t^{N'}\right)^2 > 0$, which is always satisfied if $t^{A'} \neq t^{N'}$.

In general the export supply function can have increasing or decreasing elasticity, so this is ultimately an empirical question. It is interesting to relate this analysis with a central result of the TOT theory of trade agreements, highlighted by Bagwell and Staiger (1999) and other papers by the same authors, namely that a mutually beneficial TA always expands trade relative to the noncooperative equilibrium. Recast in our framework, the analog of Bagwell and Staiger's result is that the *mean motive* for a TA has an unambiguous expanding impact on trade. In contrast, the *uncertainty motive* for a TA may impact trade volume in either direction.[33]

There is a special but interesting case where the model yields a more definite prediction about the impact of a decrease in trade policy uncertainty on expected trade. This is the same case we considered above when highlighting that an uncertainty-reducing motive is more likely to be present for lower-income countries. We showed above that, if preferences are Cobb-Douglas and Foreign is sufficiently specialized, then the optimal MPA reduces policy uncertainty for any $\theta < 0$. In this case, the export supply elasticity ε_x^* must be decreasing in π^* around the point of full specialization, since it is zero if the country is fully specialized (and we assumed $\varepsilon_x^* \geq 0$). As a consequence, a decrease in policy uncertainty increases expected trade. This suggests that heavily specialized countries are not only more likely to benefit from a reduction in policy uncertainty, as we argued above, but also more likely to experience an increase in expected trade volume if policy uncertainty decreases.

IV. More general economic shocks

Thus far we have focused on shocks of the political-economy kind, which affect Foreign welfare only through Home's tariff t. We now extend the analysis to the case of more general economic shocks, allowing λ to affect Foreign welfare not just through the policy but also directly; conventional demand or supply shocks in Home and/or in Foreign in general will have this feature. This extension is important for two reasons. First, empirically there is evidence that trade policy responds to a variety of economic shocks such as aggregate downturns (see Bown and Crowley, 2013). Second,

[33] One can also ask how the optimal MPA affects trade volatility. It is easy to show that in the "neutral" case of constant export supply elasticity, an MPA that reduces policy uncertainty also reduces trade volatility. Thus there is a tendency for the optimal MPA to impact policy uncertainty and trade (volume and value) uncertainty in the same direction. But if the export supply elasticity is not constant, the impact of a change in policy uncertainty on trade volatility is ambiguous.

economic shocks may magnify or dampen the impact of Home's trade protection on Foreign, that is, they may have a policy-externality-shifting effect, in addition to the policy-risk-preference effect.

To apply the condition derived in the reduced-form analysis of section II, start by recalling that Foreign's terms-of-trade are given (in logarithmic form) by $\pi^*(t,\lambda) = \pi(\lambda) - t - \gamma$. This notation emphasizes that the shock may affect Foreign's TOT, holding the policy t constant, through Home's domestic price; this will be the case if the domestic shock affects economic conditions at Home. In addition to affecting Foreign welfare through the TOT channel just highlighted, the shock may also affect Foreign welfare *directly* (that is, holding the TOT constant); this will be the case for example if λ represents a global demand or supply shock.

We extend our notation to reflect the more general nature of the shock. To this end, we write Foreign welfare as a function of TOT and the shock as $u^*(\pi^*(\cdot),\lambda)$. Recalling that the Foreign government maximizes national welfare, we can then write $G^*(t,\lambda) = u^*(\pi(\lambda) - t - \gamma, \lambda)$.

Recall from section II that there is an uncertainty-reducing motive for a TA if $G_{tt}^{*N} \cdot \frac{dt^N}{d\lambda} + G_{t\lambda}^{*N} < 0$, and recall our interpretation of the term $G_{tt}^{*N} \cdot \frac{dt^N}{d\lambda}$ as capturing the effect of policy-risk preference, while we interpreted the term $G_{t\lambda}^{*N}$ as capturing a policy-externality-shifting effect.

In what follows it is convenient to interpret λ as the log of the underlying shock, so that $\varepsilon_\lambda^\tau \equiv t^{N\,\prime}(\lambda)$ can be interpreted as the elasticity of the tariff factor with respect to the shock.

Using $G_t^{*N} = v^{*\theta}\Omega^*$, plugging in the expression (III.2) for G_{tt}^{*N} and simplifying, we find that there is an uncertainty-reducing motive for a TA if

$$\text{(IV.1)} \qquad (\theta\Omega^* + \varepsilon_x^* + D^*)\left(\varepsilon_\lambda^\tau - \varepsilon_\lambda^\pi\right) - \frac{\partial \ln\left(v^{*\theta}\Omega^*\right)}{\partial\lambda} < 0,$$

where $\frac{\partial \ln\left(v^{*\theta}\Omega^*\right)}{\partial\lambda}$ denotes the elasticity of adjusted openness with respect to the shock holding π^* constant, $\varepsilon_\lambda^\pi \equiv \pi'(\lambda)$ is the elasticity of Home's domestic price with respect to the shock, and (IV.1) is evaluated at the noncooperative tariff.

To interpret (IV.1), start by recalling that the sign of Foreign's preference for trade policy risk is given by the sign of $(\theta\Omega^* + \varepsilon_x^* + D^*)$. Thus the term $(\theta\Omega^* + \varepsilon_x^* + D^*)\,\varepsilon_\lambda^\tau$ in (IV.1) is related to the policy-risk preference effect. This term is analogous to the case of political-economy shocks considered in the previous section.

The new feature with more general shocks is the presence of a policy-externality-shifting effect. Recall our discussion above of the two possible channels through which λ can affect Foreign welfare holding t constant. Similarly, λ can affect the marginal international externality through two possible channels: the term $(\theta\Omega^* + \varepsilon_x^* + D^*)\,\varepsilon_\lambda^\pi$ in (IV.1) captures the impact of λ on the policy externality through Home's domestic price π, and the term $\frac{\partial \ln\left(v^{*\theta}\Omega^*\right)}{\partial\lambda}$ captures the direct impact of λ on the policy externality holding the TOT, π^*, constant.

First focus on the case in which the shock λ is *importer specific*, in the sense that it originates in the Home country and affects Foreign welfare only through the TOT. In this case only the first of the two channels highlighted above is operative, so $\frac{\partial \ln\left(v^{*\theta}\Omega^*\right)}{\partial\lambda} = 0$ and condition (IV.1) boils down to $(\theta\Omega^* + \varepsilon_x^* + D^*)\,(\varepsilon_\lambda^\tau - \varepsilon_\lambda^\pi) < 0$. To highlight the implications of this type of shock, suppose that Foreign is averse to TOT risk (or equivalently to trade-policy risk), that is $\theta\Omega^* + \varepsilon_x^* + D^* < 0$. Note that the total impact of λ on TOT is given by $\frac{d\pi^*}{d\lambda} = \varepsilon_\lambda^\pi - \varepsilon_\lambda^\tau$, so there are two different sources of TOT risk: a "policy" risk (captured by $\varepsilon_\lambda^\tau > 0$) and an "economic" risk (captured by ε_λ^π). Without economic risk (e.g. in the case of a pure political-economy shock), a mean preserving compression in t clearly reduces TOT risk. And the same is true whenever policy risk is not offset by economic risk, so that $\frac{d\pi^*}{d\lambda} < 0$. But if the economic risk offsets the policy risk (ε_λ^π is positive and dominates ε_λ^τ), then TOT risk is reduced by *increasing* policy risk, so in this case the optimal MPA will increase policy risk.[34]

Next focus on the case in which the shock λ is *global*, in the sense that it affects domestic conditions in both countries (or equivalently, suppose that the two countries experience perfectly correlated domestic shocks). In this case both channels of the policy-externality-shifting effect that we described above will be operative. The second effect (through $\frac{\partial \ln\left(v^{*\theta}\Omega^*\right)}{\partial\lambda}$) can be interpreted as follows: if shocks that increase the noncooperative tariff also increase the adjusted degree of openness for a fixed tariff, this strengthens the uncertainty-reducing motive.

[34] In the case of importer-specific shocks we can show a further result: under a regularity condition that we specify below, the optimal MPA reduces *terms-of-trade* risk if Foreign is averse to TOT risk (or equivalently to trade-policy risk), that is if $\theta\Omega^* + \varepsilon_x^* + D^* < 0$. Thus, the impact of the optimal MPA on TOT risk is determined solely by the Foreign country's preference for TOT/policy risk, and follows the same intuitive pattern as in the case of political economy shocks. The regularity assumption we need is the following: if we define Home's choice variable as π^* rather than t (which is clearly equivalent), we need Home's noncooperative choice of π^* to be monotonic in λ, which is ensured if $\frac{d^2 G}{d\pi^* d\lambda}$ does not change sign over the relevant range of (π^*, λ).

It is worth emphasizing that, unlike in the case of political-economy shocks considered in the previous section, here the direction of the uncertainty motive for a TA may go in a different direction than Foreign's preference for policy risk. So, for example, it is *possible* that even if individuals are risk-neutral ($\theta \to 0$) and hence the Foreign country is policy-risk loving, there may be an uncertainty-reducing motive for a TA.

The sign of the externality-shifting effect in general depends on the exact nature of the shock and of the economic structure, but we highlight an interesting case in which the externality-shifting effect pushes towards an uncertainty-reducing motive. Suppose that λ is a global productivity shock that strengthens comparative advantage, so that Foreign's openness Ω^* is higher (for given TOT) when λ is higher. Further suppose that Home's noncooperative tariff t^N increases with trade volume; this is compatible with our model if TOT manipulation motives are important for Home's choice of tariff. In this case t^N is increasing in λ, as assumed in our model. Then, if the effect of the shock via $v^{*\theta}$ is not too strong, the sign of $\frac{\partial \left(v^{*\theta}\Omega^*\right)}{\partial \lambda}$ will be positive, thus contributing towards an uncertainty-reducing motive.

V. A sufficient statistic for the uncertainty motive

If one is willing to assume that the model is true, one can in principle use the model to determine empirically the direction of the uncertainty motive for a TA between two countries and quantify the relative gains from regulating policy uncertainty. In this section we illustrate with a simple example how this could be done with actual data.

As we observed above, there is an uncertainty-reducing motive for a TA if the (negative) international externality from the tariff at the noncooperative equilibrium is stronger when λ is higher, that is if $\frac{d}{d\lambda}(-v^{*\theta}\Omega^*)^N < 0$. Since $t^N(\lambda)$ is increasing, this condition can be equivalently written as

$$\text{(V.1)} \qquad \frac{d(v^{*\theta}\Omega^*)^N}{dt^N} > 0.$$

Note that this condition is valid not only in the case of political-economy shocks considered in section III, but also in the case of more general economic shocks considered in section IV.

In principle, condition (V.1) can be implemented empirically, if one has data on a small country facing non-cooperative tariffs from the rest of the world. Suppose one has information on this country's openness (Ω^*), real income per capita (v^*) and estimates of θ to construct a measure of the adjusted degree of openness, as well as the average tariff faced by this country (t^N). Our model then suggests that, *if the adjusted measure of openness co-varies with the tariff, then there is an uncertainty-reducing motive for a TA.*

This sufficient-statistic approach can also be used to approximate the relative gains from regulating policy uncertainty. Applying formula (II.7), we can write

$$\frac{\tilde{V}^{MPA}}{\tilde{V}^{UPA}} = \left| \frac{d \ln(v^{*\theta}\Omega^*)^N}{dt^N} \right| \cdot \frac{\sigma^2_{t^N}}{\bar{t}^N}$$

This suggests quantifying $\tilde{V}^{MPA}/\tilde{V}^{UPA}$ by taking a measure of correlation between $\ln(v^{*\theta}\Omega^*)^N$ and t^N, for example the estimated coefficient of a simple OLS regression, and multiplying it by $\sigma^2_{t^N}/\bar{t}^N$. In what follows we will adopt this approach and quantify $\tilde{V}^{MPA}/\tilde{V}^{UPA}$ as $\left|\beta^{ols}\right| \cdot \frac{\sigma^2_{t^N}}{\bar{t}^N} = \frac{\left|Cov\left(\ln(v^{*\theta}\Omega^*)^N, t^N\right)\right|}{\bar{t}^N}$, where β^{ols} is the estimated OLS coefficient.[35]

We now illustrate how this approach can be implemented with actual data by focusing on a simple empirical example, namely the trade relationship between US and Cuba in the period before 1934. As already mentioned in the introduction, this was a period of non-cooperative trade relations, which ended with the Reciprocal Trade Agreement Act (RTAA). The first agreement signed by the US under the RTAA was the agreement with Cuba in 1934. This, together with the fact that Cuba was a small open country (its export share of GDP in this period was on average 0.32) that exported mostly to the US, makes these countries a good fit to illustrate our approach.

Our model is static in nature, but it seems natural to use the time variation in noncooperative tariffs and adjusted openness to measure their covariation. We focus on the annual US average tariff prior to 1934. More specifically, we use $t = \ln(1 + \tau)$, where τ is the US import-weighted average tariff starting in 1867 calculated by Irwin (2007). Figure 2 plots t from 1867 to 1960, showing considerable variation prior to

[35] As we discussed in section A, the ratio $\tilde{V}^{MPA}/\tilde{V}^{UPA}$ involves gains from small policy changes, rather than gains from optimal policy changes, so it is not the ideal measure, but we see no reason to believe that $\tilde{V}^{MPA}/\tilde{V}^{UPA}$ would systematically overstate or understate the relative gains from optimal policy changes.

1934.[36] We use data available for Cuba on openness and income per capita in the period 1903-1933 to calculate a measure of adjusted openness at alternative levels of risk aversion.[37]

The first point we note is that, if citizens were income-risk neutral ($\theta = 0$), to evaluate the direction of the uncertainty motive we would only need to look at the sign of the covariance between Cuban openness and the US tariff. We find this covariance to be negative, which is plausible, since higher US tariffs tend to reduce the Cuban share of exports in GDP, and is consistent with the model, recalling the result that if $\theta = 0$ there should be an uncertainty-increasing motive for a TA (see section III).[38]

If, as is more reasonable, citizens are risk averse, then we need to consider the covariance between Cuba's *adjusted* openness and the US tariff. We compute this covariance at alternative levels of θ, and find that it is positive for $\theta < -1.1$. So our analysis indicates that there is an uncertainty-reducing motive for a TA for $\theta < -1.1$. We do not have estimates of risk aversion for Cuba, but we note that Kimbal, Sahm, and Shapiro (2008) estimate CRRA coefficients for US households (by using their preferences over different gambles), finding that about 90 percent of the distribution lies below -1.5.

We obtain a similar result if instead of the aggregate US tariff we use the US tariff on Cuban sugar. The latter may be a better proxy of the US trade barriers that affected Cuba directly, since Cuban exports of sugar to the US accounted for 25-30 percent of Cuban national income (Dye, 2005, p. 193). Using the sugar tariff, we find

[36] Part of this variation is simply a downward trend, but there is also considerable variation around the trend. This trend is probably due to the fact that the revenue motives for imposing tariffs (which were arguably important before the civil war) declined over time for various reasons, including the introduction of the income tax in 1916. Another part of the variation is caused by price changes since the US had many specific tariffs. However, statutory rates also oscillated considerably prior to 1934 depending on whether Congress was controlled by Republicans (protectionist) or Democrats. The RTAA lowered the ability of Congress to engage in such policy reversals.

[37] The start date is dictated by income data availability from the Montevideo-Oxford Latin American Economic History Database, available at <http://oxlad.qeh.ox.ac.uk/results.php>. We note that 1903 also coincided with an initial US-Cuba trade agreement whereby the US granted a 20% preferential reduction to Cuban sugar and tobacco. However, as Cuba scholars such as Dye and Sicotte (1999) point out, there was no legal commitment to those lower tariffs so the "regime was not risk-free – exporters in both countries faced the possibility that tariff modifications could reduce or even eliminate the benefits conveyed by the treaty" (p. 22). This was in fact what happened starting in 1921 when the US increased tariffs on several goods including Cuban sugar. In fact, some argue that the subsequent US tariff increases in the Smoot Hawley act caused the sharp decline of Cuban Sugar exports in 1930-33 and contributed to the Revolt of 1933.

[38] We find a negative covariance between $\ln \Omega^*$ and t whether or not we control for a linear time-trend.

that there is an uncertainty-reducing motive for $\theta < -1$.

Finally, using the approach developed above, we can quantify the relative gains from regulating policy uncertainty ($\tilde{V}^{MPA}/\tilde{V}^{UPA}$) by computing the adjusted covariance measure $|Cov\left(\ln(v^{*\theta}\Omega^{*})^{N}, t^{N}\right)|/\bar{t}^{N}$ at alternative levels of θ. Table 1 reports the results of this quantification. Note from the last row of Table 1 that, even at moderate levels of risk aversion, the estimate of $\tilde{V}^{MPA}/\tilde{V}^{UPA}$ is not negligible, and it is close to 1/3 if $\theta = -5$ (the median value in the study by Kimbal, Sahm, and Shapiro, 2008), when using the US tariff on Cuban sugar.

[TABLE 1 HERE]

In sum, this section illustrates how the model can be used to evaluate the direction of the uncertainty motive for a TA between two countries and to quantify the relative gains from regulating trade-policy uncertainty. The positive correlation between US tariffs and Cuban adjusted openness at reasonable levels of risk aversion suggests that there was indeed an uncertainty-reducing motive for a TA between these two countries before 1934, and we find the relative gains from reducing policy uncertainty to be significant. It is important to emphasize, however, that this exercise is not a test of the model, but rather it assumes that the model is true and so it must be taken with a grain of caution, since the model is very stylized. The message we want to convey is that it is feasible to take our model to the data in a meaningful way, and it might be desirable to develop richer and more realistic versions of our model in order to quantify the uncertainty-related gains from TAs.

VI. Ex-ante investments

Our basic model assumes that allocation decisions occur ex post, after the shock is realized. But in reality there are a variety of production factors that cannot be flexibly shifted in response to policy and economic shocks. In this section we extend our analysis to allow for allocation decisions that must be made ex-ante, before the shock is realized, or "ex-ante investments". As we noted in the introduction, the often-heard informal arguments about the motives for TAs claim that they should increase investment and trade by reducing uncertainty. Allowing for ex-ante investments in our model seems compelling if one wants to formally examine this issue.

Recall that the standard model allows for an arbitrary number of factors that are mobile ex-post. We now assume that one of these, "capital," is mobile ex-ante but fixed ex-post.[39] We normalize the endowment of capital to one and let k^* denote the fraction of capital allocated to the export sector. To simplify the analysis we assume that all factors in the Home country are perfectly flexible so they can be allocated after the shock λ is realized. This allows us to keep the economic structure for Home in the background, as we did in the static model.

We assume the following timing: (0) The tariff schedule is selected (cooperatively or noncooperatively); (1) capital is allocated; (2) λ is realized; (3) the trade policy is implemented and markets clear.

Both in the cooperative and noncooperative scenarios, we allow the tariff schedule to be contingent on λ. Note that we keep the timing constant across the cooperative and noncooperative scenarios. The reason for this choice is to abstract from domestic-commitment motives for a TA. And of course, if we want a TA to be able to affect investment decisions by managing policy uncertainty, we need policy choices to be made before investment decisions, and this explains our choice of timing.[40]

The first step of the analysis is to extend Proposition 1 from the previous static setting to the present dynamic environment. We write Foreign welfare as $G^*(t, \lambda, k^*)$, and we continue to write Home's objective as $G(t, \lambda)$, which reflects the assumption that Foreign is a small country.[41]

In keeping with our assumption that there is no role for trade policy intervention in Foreign, we assume that capital is perfectly divisible, so that the citizens of the small country are not only identical ex-ante, but also ex-post, and thus there is no redistribution motive for a tariff. This in turn implies that, given Home's (cooperative or noncooperative) tariff schedule $t(\lambda)$, capital in Foreign is efficiently allocated, and

[39] We could allow for a higher number of factors that are mobile ex ante but fixed ex post, but the notation would get more cumbersome. And of course, the model also allows for factors that are fully fixed (immobile both ex ante and ex post).

[40] While the assumption is made to provide a clean thought experiment, we note that in some cases countries are able to unilaterally choose contingent protection programs in ways that represent long-term commitments. For example the U.S. and the E.U. have contingent protection laws that apply in the absence of trade agreements.

[41] If Home's objective G is some weighted social welfare function, then for a given Home tariff t the level of k^* can affect G only through the Home country's terms of trade π, but since Foreign is small π is not affected by k^*. On the other hand, k^* can in general affect the noncooperative tariff t^N, for example because it can affect the Foreign country's export supply elasticity. In our notation we suppress the dependence of t^N on k^*, as this should not cause any confusion.

hence k^* maximizes $EG^*(t(\lambda), \lambda, k^*)$.[42] To simplify the arguments below, we assume that G^* is strictly concave in k^*.

As in the previous static setting, we characterize the optimal MPA, that is the tariff schedule that maximizes expected joint welfare subject to the constraint $Et(\lambda) = Et^N(\lambda)$.

We now argue that Proposition 1 extends to this setting, in the sense that we only need to determine the sign of $\frac{d}{d\lambda}G_t^*(t^N(\lambda), \lambda, k^*)$ to know if there is an uncertainty-reducing role for a TA. The following local argument provides some intuition for the result. Starting at $t^N(\lambda)$, a small mean-preserving compression has no first order effect on EG since this objective is maximized by $t^N(\lambda)$. Therefore, the new schedule will only increase EG^W if it increases EG^*. Since, as noted above, k^* maximizes $EG^*(t(\lambda), \lambda, k^*)$, this policy change has no first-order effect on EG^* via k^*. So any impact of the policy change on EG^* must be due to the "static" effect, i.e. to $\frac{d}{d\lambda}G_t^{*N} \neq 0$.

We now consider the full MPA program. Recalling that, for a given $t(\lambda)$, the level of k^* maximizes $EG^*(t(\lambda), \lambda, k^*)$ and has no effect on EG, then k^* maximizes $EG^W(t(\lambda), \lambda; k^*)$. Thus we can write the MPA program as if the governments were choosing k^* directly:

$$\max_{t(\lambda),k^*} EG^W(t(\lambda), \lambda, k^*) \quad \text{(VI.1)}$$
$$\text{s.t. } Et(\lambda) = Et^N(\lambda)$$

Assuming an interior optimum, we obtain the following FOCs:

$$G_t^W(t, \lambda, k^*) = \psi \text{ for all } \lambda \quad \text{(VI.2)}$$
$$Et(\lambda) = Et^N(\lambda)$$
$$EG_{k^*}^W(t(\lambda), \lambda, k^*) = 0 \quad \text{(VI.3)}$$

We can now apply an argument similar to the static model, using the first two of the FOC above. The only difference is that the derivative $\frac{d}{d\lambda}G_t^{*N}$ is evaluated at the optimal level of k^*, but as long as the sign of this derivative does not change with k^*,

[42] If capital is divisible, all citizens have identical incomes ex-post, and as a consequence there is no idiosyncratic risk, which implies that the competitive allocation is efficient, conditional on Home's trade policy. Note that there is aggregate risk in this economy, but it cannot be diversified away (since there are no international insurance markets in our model).

Proposition 1 extends to this setting. In Appendix we prove the following:

Proposition 4. *If $\frac{d}{d\lambda}G_t^*(t^N(\lambda), \lambda, k^*) < 0$ (> 0) for all (k^*, λ), then there is an uncertainty-reducing (-increasing) motive for a TA.*

Proposition 4 highlights that the uncertainty motive for the TA is driven by the static effect, i.e. the impact of the shock on the policy externality conditional on the capital level. In a broad sense, we can interpret this result as indicating that the presence of ex-ante investments does not generate a *separate* uncertainty motive for a TA.

This conclusion, as we highlighted, relies on the competitive allocation of capital being socially efficient given Home's trade policy, which is ensured in our setting by the assumption of perfectly divisible capital. While this assumption is somewhat restrictive, we note that the same result would obtain in a setting where capital is not divisible, provided that an efficient domestic insurance market is present, or alternatively that the government can use an entry subsidy/tax to control the allocation of capital.[43]

Of course one could consider reasonable alternative scenarios where capital allocation is not efficient, and in such scenarios there could be an "investment motive" for an MPA, or in other words, there could be scope for a TA to "correct" the capital allocation through changes in policy uncertainty, but we note that this would be a second-best argument for a TA, as the first-best way to address such inefficiency would be the use of more targeted policies.

Given that the condition for an uncertainty-reducing motive for a TA is similar as in the static model, the results of the previous sections all extend to the present setting, with the only difference that the relevant expressions are evaluated at a given capital allocation. Moreover, the expressions for the approximate values of an MPA and a UPA are also unchanged, since there is no first order effect on Foreign welfare due to capital re-allocation. But even if there is no separate "investment motive" for an MPA, such an agreement in general does affect equilibrium investment levels relative to the noncooperative equilibrium, as we show next.

[43] If capital is indivisible, so that each citizen must choose ex-ante whether to allocate her capital to the export sector or the import-competing sector, then ex-post agents fare differently in different states of the world. In this situation, the competitive equilibrium is efficient (given Home's trade policy) only if a domestic insurance market is present, or if the government can use policies to correct the allocation of capital, such as an entry subsidy/tax.

A. Impact of policy uncertainty on investment and trade

We start by asking how the optimal MPA affects ex-ante investments. We focus on the case in which $\frac{d}{d\lambda}G_t^{*N} < 0$, so that the optimal MPA reduces policy risk. To simplify the exposition we assume that the trade pattern does not switch as k^* changes, that is, Foreign exports the nonnumeraire good for all $k^* \geq 0$. Also, for simplicity we focus here on the case of political economy shocks, as in the basic model of section III.

Recall that efficient capital allocation implies $\frac{\partial EG^*}{\partial k^*} = 0$. By standard results (Rotschild and Stiglitz, 1971), the equilibrium k^* increases as a result of a mean-preserving compression in t if $\frac{\partial}{\partial k^*}G_{tt}^*(t, k^*) < 0$ for all t in its support. Thus the effect depends on the impact of k^* on Foreign's policy-risk preference. In general this effect can go in either direction, but we now highlight a set of sufficient conditions under which it is negative.[44]

Note that the result of Proposition 3 extends directly to this dynamic setting, in the sense that the expression for G_{tt}^* is just the same as in (III.2), provided its various components are re-interpreted as *conditional on the capital allocation* k^*. Subject to this re-interpretation, we have

$$\frac{\partial}{\partial k^*}G_{tt}^*(t, k^*) = \frac{\partial}{\partial k^*}\left[v^{*\theta}\Omega^*\left(\theta\Omega^* + \varepsilon_x^* + D^*\right)\right] \tag{VI.4}$$

In Appendix we prove that, if θ is sufficiently negative and the support of λ sufficiently small, then $\frac{\partial}{\partial k^*}G_{tt}^*(t, k^*) < 0$ for all t in its support, which leads to the following:

Proposition 5. *Suppose λ is a political economy shock. If there is sufficient income risk aversion and the support of λ is sufficiently small, then the optimal MPA increases investment in the export sector.*

Broadly interpreted, this proposition suggests that under the condition that generates an uncertainty-reducing motive for a TA, namely a strong degree of income-risk

[44] The general ambiguity of the impact of mean-preserving changes in prices on investment decisions is well known. In the literature this ambiguity is resolved in different ways, e.g. assuming decreasing absolute risk aversion, positing a specific shock distribution, restricting the economic environment or, as we do, considering cases with small uncertainty. But we emphasize that our result is novel: we are not aware of any existing result that expresses a similar set of sufficient conditions for a similar economic environment. We also note that we could prove the result under the alternative assumption that the probability mass is sufficiently concentrated, rather than the support being sufficiently small, but in this case the notation and the analysis would be more cumbersome.

aversion, the optimal MPA leads to higher investment in the export sector, provided the underlying uncertainty in the environment is small enough. We also note that the same result would hold if we replaced the condition that θ is sufficiently negative with the alternative condition that the export supply elasticity ε_x^* is sufficiently close to constant, as we show in Appendix.

Finally we examine the impact of the optimal MPA on expected trade volume in the presence of ex-ante investments.

Recall first that, in the absence of ex-ante investment, if the MPA reduces policy uncertainty, expected trade increases if and only if the export supply elasticity $\varepsilon_x^*(\pi^*)$ is decreasing in π^*. In the presence of ex-ante investment, we can write trade volume as $x^*(\pi^*, k^*)$, thus the MPA increases expected log trade if and only if the following is positive

$$\begin{aligned} & \int \ln x^*\left(\pi^*, k^{*MPA}\right) dF_{MPA}^k(\pi^*) - \int \ln x^*\left(\pi^*, k^{*N}\right) dF_N^k(\pi^*) \\ = & \int \ln x^*\left(\pi^*, k^{*N}\right) d\left(F_{MPA}^k(\pi^*) - F_N^k(\pi^*)\right) + \int \ln \frac{x^*\left(\pi^*, k^{*MPA}\right)}{x^*\left(\pi^*, k^{*N}\right)} dF_{MPA}^k(\pi^*) \end{aligned}$$

where k^{*MPA} and k^{*N} are respectively the equilibrium capital levels at the optimal MPA and at the noncooperative equilibrium, and F_N^k and F_{MPA}^k are the respective distributions of π^*. The first term in the expression above is analogous to the one in the static model, so it depends on whether $\varepsilon_x^*\left(\pi^*, k^{*N}\right) \equiv \partial \ln x^*\left(\pi^*, k^{*N}\right)/\partial\pi^*$ is increasing or decreasing in π^*. The second term captures the expected growth in exports due to the change in investment. If k^* increases, this effect will be positive if the support of the shock is sufficiently small and the economy is not completely specialized.[45]

Summarizing the discussion above, if risk aversion is sufficiently strong and uncertainty is sufficiently small, the optimal MPA reduces uncertainty in trade policy

[45] To see this, note that $\frac{\partial x^*(\pi^*,k^*)}{\partial k^*} = \frac{\partial(q^*-c^*)}{\partial k^*} = \frac{\partial q^*}{\partial k^*} - \frac{\partial c^*}{\partial R^*} \cdot \frac{\partial R^*}{\partial k^*}$, where $\frac{\partial R^*}{\partial k^*}$ is the ex-post differential in the rate of return to capital across sectors. This differential is zero in expectation under risk neutrality, while it can differ from zero with risk aversion, but if the shock has small support it is close to zero at the optimal ex-ante allocation. Thus if the support of λ is sufficiently small then $\frac{\partial x^*}{\partial k^*} > 0$, provided that $\frac{\partial q^*}{\partial k^*} > 0$, which is the case if the economy is not completely specialized.

One may also ask how an MPA affects the volatility of trade flows. When ε_x^* is not constant, this impact is ambiguous, but it is direct to show that in the "neutral" case where ε_x^* is constant, an MPA that decreases trade policy uncertainty decreases uncertainty in trade volume, i.e. $\ln x^*\left(\pi^{*N}\right)$ is a MPS of $\ln x^*\left(\pi^{*MPA}\right)$.

and increases investment in the export sector. Moreover, under these conditions, expected trade increases provided the export supply elasticity does not increase too rapidly with the price.

We conclude this section with a final point regarding the statement made by the WTO that one of its key goals is to reduce policy uncertainty for the purposes of increasing investment in export sectors. Our analysis suggests that, even though under some conditions a reduction in policy uncertainty does lead to more investment in the export sector, this by itself does not imply a first-order welfare increase: if capital markets are efficient, the only first-order welfare change from a (small) reduction in policy uncertainty is of a "static" nature, that is, it comes from the correction of the international policy-risk externality, conditional on the initial allocation of capital.

VII. Two policy-active countries

In this section we extend our analysis by considering a setting with two policy-active countries. We focus on the reduced-form framework of section II and abstract from ex-ante investments for simplicity.

We represent the reduced-form payoff functions as $G(t, t^*, \lambda)$ and $G^*(t^*, t, \lambda^*)$, where t is Home's policy and t^* is Foreign's policy. For tractability, we assume that countries are mirror-image symmetric, and we continue to assume a single dimension of uncertainty, that is $\lambda^* = \lambda$; the interpretation is that there is a global shock that affects the two countries symmetrically, or equivalently, two domestic shocks that are perfectly correlated. We assume that each payoff function is concave in its first argument ($G_{tt} < 0, G^*_{t^*t^*} < 0$), that the single-crossing property is satisfied ($G_{t\lambda} > 0, G^*_{t^*\lambda} > 0$) and that reaction functions are stable ($|G_{tt}| > G_{tt^*}$).

We denote the common payoff given a symmetric tariff t as $\tilde{G}(t, \lambda) \equiv G(t, t, \lambda)$. We assume that $\tilde{G}$ is concave in t and satisfies the single crossing property ($\tilde{G}_{t\lambda} > 0$).

Given that countries are symmetric, we look for a symmetric noncooperative equilibrium tariff, which is implicitly defined by the following FOC:

$$G_t(t^N, t^N, \lambda) = 0.$$

Let $t^N(\lambda)$ denote the noncooperative tariff schedule. Given our assumptions, $t^N(\lambda)$

is increasing, as can be verified by implicitly differentiating the FOC:

$$\frac{dt^N}{d\lambda} = \frac{G^N_{t\lambda}}{-(G^N_{tt} + G^N_{tt^*})} > 0$$

where the numerator is positive by the single crossing property and the denominator is positive by the stability assumption.

Given the symmetry of the problem, it is natural to focus on the optimal symmetric MPA,[46] which is given by:

$$\text{(VII.1)} \qquad t^{MPA}(\lambda) = \arg\max_{t(\lambda)} E\tilde{G}(t(\lambda), \lambda) \text{ s.t. } Et(\lambda) = Et^N(\lambda).$$

We can write the Lagrangian for this problem as

$$\text{(VII.2)} \qquad L = \int [\tilde{G}(t, \lambda) + \psi \left(t^N(\lambda) - t(\lambda)\right)] dF(\lambda)$$

Maximizing this Lagrangian pointwise yields the FOCs

$$\begin{aligned} \tilde{G}_t(t(\lambda), \lambda) &= \psi \text{ for all } \lambda \\ Et(\lambda) &= Et^N(\lambda) \end{aligned}$$

We can then prove the following:

Proposition 6. *If* $(G^{*N}_{tt} + G^{*N}_{tt^*}) \cdot \frac{dt^N}{d\lambda} + G^{*N}_{t\lambda} < 0$ (> 0) *for all* λ *then there is an uncertainty-reducing (-increasing) motive for a TA. If* $(G^{*N}_{tt} + G^{*N}_{tt^*}) \cdot \frac{dt^N}{d\lambda} + G^{*N}_{t\lambda} = 0$ *for all* λ *then there is no uncertainty motive for a TA.*

We can now contrast the result of Proposition 6 with the corresponding result for the small-large country setting. The general condition for an uncertainty-reducing motive, $\frac{d}{d\lambda} G^{*N}_t < 0$, is similar as in the small-large country setting, but in the large-large country setting this expression includes an additional term, namely $G^{*N}_{tt^*}$. We label this the "strategic interaction" effect, which is positive if tariffs are strategic complements and negative if they are strategic substitutes. Thus an interesting new insight that emerges is that the strategic-interaction effect works in favor of the

[46] Given the concavity of the payoff functions, we conjecture that the global maximum is indeed symmetric.

uncertainty-reducing motive if tariffs are strategic substitutes, and vice-versa if tariffs are strategic complements. Whether tariffs are strategic substitutes or complements depends on the specifics of the trade structure (see for example Syropoulos, 2002), so the direction of this effect is ultimately an empirical question.

Note also that, while the other terms are similar as in the small-large country setting, they will reflect additional effects when one applies the general formula to a specific trade structure. In particular, the policy-risk-preference effect G_{tt}^{*N} and the externality-shifting effect $G_{t\lambda}^{*N}$ will include tariff-revenue and pass-through elasticity effects that were absent in the small-large country setting.

Finally, it can be shown that the expressions derived in section A for the gains from regulating policy uncertainty and policy mean extend directly to the large-large country setting considered in this section.

VIII. Conclusion

The objective of this paper is to conduct a rigorous examination of the often-heard informal argument that an important motive for TAs is to reduce uncertainty in trade barriers. Focusing on a standard competitive trade model with political/economic shocks, we find that if citizens are risk neutral there tends to be an uncertainty-increasing motive for a TA. If citizens are risk-averse, an uncertainty-reducing motive for a TA is more likely to be present, other things equal, when the economy is more open, the export supply elasticity is lower, the economy is more specialized, and citizens are more risk-averse. The model suggests that, as the world becomes more integrated, the gains from decreasing trade-policy uncertainty should tend to become more important relative to the gains from reducing the levels of trade barriers. Furthermore, governments have more to gain by joining a TA when the trading environment is more uncertain. We develop a simple "sufficient statistic" approach to determine the direction of the uncertainty motive for a TA and quantify the associated gains, and illustrate how it can be taken to the data. Finally, we examine how the uncertainty motive for a TA is affected by the presence of ex-ante investments, and examine conditions under which an uncertainty-reducing TA will increase investment in the export sector and raise expected trade volume.

There are several potentially interesting avenues for future research. First, in this paper we have abstracted from contracting frictions. As mentioned in the introduction, we believe this is a natural first step given that our main focus is the potential gains from regulating policy uncertainty, but it would be interesting to examine how results would change in the presence of contracting frictions. Second, it would be desirable to examine the potential uncertainty-managing role of trade agreements in settings where the underlying reason for the agreement is not the classic TOT externality: in particular, one might consider settings in which agreements are motivated by the governments' need for domestic commitment, or by the presence of non-TOT international externalities. Finally, a challenging but potentially fruitful direction of research would be to develop a richer version of our model with the objective of taking it to a comprehensive dataset: this would probably require, among other things, allowing for multiple countries, multiple goods and imperfectly correlated shocks across countries.

IX. References

Amador, Manuel, and Kyle Bagwell. 2013."The Theory of Optimal Delegation with an Application to Tariff Caps." *Econometrica* 81(4): 1541–99.

Anderson, James E., and John G. Riley. 1976. "International Trade with Fluctuating Prices." *International Economic Review* 17(1): 76–97.

Bagwell, Kyle, and Robert W. Staiger. 2012. "Profit Shifting and Trade Agreements in Imperfectly Competitive Markets." *International Economic Review* 53(4): 1067–1104.

Staiger, Robert W., and Kyle Bagwell. 1999. "An Economic Theory of GATT." *American Economic Review* 89(1): 215–248.

Beshkar, Mostafa, and Eric W. Bond. 2012. "Cap and Escape in Trade Agreements." http://pages.iu.edu/ mbeshkar/papers/Cap-Escape.pdf.

Blattman, Christopher, Jason Hwang, and Jeffrey G. Williamson. 2007. "Winners and losers in the commodity lottery: The impact of terms of trade growth and volatility in the Periphery 1870–1939."*Journal of Economic Development* 82: 156–179.

Bown, Chad P., ed. 2011. *The Great Recession and Import Protection: The Role of Temporary Trade Barriers.* London: Centre for Economic Policy Research and the World Bank.

Bown, Chad P., and Meredith A. Crowley. 2013. "Import Protection, Business Cycles, and Exchange Rates: Evidence from the Great Recession." *Journal of International Economics* 90(1): 50-64

Cadot, O., Olarreaga, M. and Tschopp, J. 2010. "Do Trade Agreements Reduce the Volatility of Agricultural Distortions?"In *The Political Economy of Agricultural Trade Distortions*, edited by Kym Anderson. Cambridge: Cambridge University Press.

Dye, Alan D. 2005. "Cuba and the Origins of the US Sugar Quota." *Revista de Indias* LXV: 193–218.

Dye, Alan, and Richard Sicotte. 1999. "US-Cuban Trade Cooperation and its Unraveling." *Business and Economic History* 28(2): 19–31.

Eaton, Jonathan. 1979. "The Allocation of Resources in an Open Economy with Uncertain Terms of Trade." *International Economic Review* 20(2): 391–403.

Eaton, Jonathan, and Gene M. Grossman. 1985. "Tariffs as Insurance: Optimal Commercial Policy when Domestic Markets are Incomplete." *Canadian Journal of Economics* 18(2): 258–72.

Flemming, John S., Stephen J. Turnovsky, and Murray C. Kemp. 1977. "On the Choice of Numeraire and Certainty Price in General Equilibrium Models of Price Uncertainty." *Review of Economic Studies* 44(3): 573–83.

Grossman, Gene M., and Elhanan Helpman. 1995. “Trade Wars and Trade Talks.” *Journal of Political Economy* 103(4): 675–708.

Handley, Kyle. 2014. “Exporting Under Trade Policy Uncertainty: Theory and Evidence.” *Journal of International Economics* 94(1): 50–66.

Handley, Kyle, and Limão, Nuno. 2012. “Trade and Investment under Policy Uncertainty: Theory and Firm Evidence.”National Bureau of Economic Research (NBER) Working Paper 17790.

Horn, Henrik, Giovanni Maggi, and Robert W. Staiger. 2010. “Trade Agreements as Endogenously Incomplete Contracts.” *American Economic Review* 100(1): 394–419.

Irwin, Douglas A. 2007.“Trade Restrictiveness and Deadweight Losses from U.S. Tariffs, 1859–1961.”National Bureau of Economic Research (NBER) Working Paper 13450.

Irwin, Douglas A. 1998. “From Smoot-Hawley to Reciprocal Trade Agreements: Changing the Course of U.S. Trade Policy in the 1930s.” In *The Defining Moment: The Great Depression and the American Economy in the Twentieth Century*, edited by Bordo, Michael D., Claudia Goldin, and Eugene N. White, 325–352. Chicago: University of Chicago Press.

Kimball, Miles S., Claudia R. Sahm, and Matthew D. Shapiro. 2008. “Imputing Risk Tolerance from Survey Responses.” *Journal of the American Statistical Association* 103(483): 1028–1038.

Limão, Nuno, and Tovar, Patricia. 2011. “Policy Choice: Theory and Evidence from Commitment via International Trade Agreements.” *Journal of International Economics* 85(2): 186–205.

Maggi, Giovanni, and Andres Rodriguez-Clare. 1998. “The value of Trade Agreements in the Presence of Political Pressures.” *Journal of Political Economy* 106(3): 574–601.

Mansfield, Edward D., and Eric Reinhardt. 2008. “International Institutions and the volatility of International Trade.”*International Organization* 62(04): 621–652.

McLaren, John. 1997. “Size, Sunk Costs, and Judge Bowker’s Objection to Free Trade.” *American Economic Review* 87(3): 400–420.

Mrazova, Monika. 2011. “Trade Agreements when Profits Matter.”http://www.monikamrazova.com/Trad

Newbery, David MG, and Joseph E. Stiglitz, eds. 1981. *The Theory of Commodity Price Stabilization: A Study in the Economics of Risk.* Oxford: Clarendon Press.

Ossa, Ralph. 2011. “A New Trade Theory of GATT/WTO Negotiations.” *Journal of Political Economy* 119(1): 122–152.

Rose, Andrew K. 2004. “Do WTO Members Have More Liberal Trade Policy?” *Journal of International Economics* 63(2): 209–235.

Rothschild, Michael, and Joseph E. Stiglitz. 1970. “Increasing Risk I: A Definition.” *Journal of Economic Theory* 2(3): 225–243.

Rothschild, Michael, and Joseph E. Stiglitz. 1971. "Increasing Risk II: Its Economic Consequences." *Journal of Economic Theory* 3(1): 66–84.

Schularick, Moritz, and Solomos Solomou. 2011. "Tariffs and economic growth in the first era of globalization." *Journal of Economic Growth* 16(1): 33–70.

Staiger, Robert W., and Guido Tabellini. 1987. "Discretionary Trade Policy and Excessive Protection." *American Economic Review* 77(5): 823–37.

Syropoulos, Constantinos. 2002. "Optimum Tariffs and Retaliation Revisited: How Country Size Matters." *The Review of Economic Studies* 69(3): 707–727.

Tokarick, Stephen. 2014. "A Method for Calculating Export Supply and Import Demand Elasticities." *The Journal of International Trade & Economic Development* (ahead-of-print): 1–29.

Young, Leslie, and James E. Anderson. 1982. "Risk Aversion and Optimal Trade Restrictions." *The Review of Economic Studies* 49(2): 291–305.

X. Appendix

Proof of Lemma 1:

We start by proving part (iii). The schedules $t^{MPA}(\lambda)$ and $t^N(\lambda)$ are clearly continuous. The mean constraint and the continuity of $t^{MPA}(\lambda)$ and $t^N(\lambda)$ ensure the existence of at least one intersection. Consider one such intersection $\hat{\lambda}$, so that $t^{MPA}(\hat{\lambda}) = t^N(\hat{\lambda})$. By the FOC, $G_t^W(t^N(\hat{\lambda}), \hat{\lambda}) = \psi$. Since $G_t(t^N(\hat{\lambda}), \hat{\lambda}) = 0$ this implies $G_t^*(t^N(\hat{\lambda}), \hat{\lambda}) = \psi$. Now if $\frac{d}{d\lambda}G_t^*(t^N(\lambda), \lambda) = 0$ then $G_t^*(t^N(\lambda), \lambda) = \psi$ for all λ, which in turn implies $G_t^W(t^N(\lambda), \lambda) = \psi$ for all λ. Therefore the schedule $t^N(\lambda)$ satisfies the FOC, hence $t^{MPA}(\lambda) = t^N(\lambda)$ for all λ.

We next prove part (i), focusing on the case $\frac{d}{d\lambda}G_t^*(t^N(\lambda), \lambda) < 0$. Again, $t^{MPA}(\lambda)$ and $t^N(\lambda)$ must intersect at least once. We now argue that $t^{MPA}(\lambda)$ can only intersect $t^N(\lambda)$ from above. This, together with continuity, will also ensure the uniqueness of the intersection.

We argue by contradiction. Suppose $t^{MPA}(\lambda)$ intersects $t^N(\lambda)$ at some point $\hat{\lambda}$ from below. Consider two values of λ on opposite sides of this intersection, $\lambda_1 < \hat{\lambda} < \lambda_2$,such that $t^{MPA}(\lambda_1) < t^N(\lambda_1)$ and $t^{MPA}(\lambda_2) > t^N(\lambda_2)$.

Recalling that $G_t(t^N(\lambda), \lambda) = 0$ and $\frac{d}{d\lambda}G_t^*(t^N(\lambda), \lambda) < 0$ for all λ, then

$$G_t^W(t^N(\lambda_2), \lambda_2) = G_t^*(t^N(\lambda_2), \lambda_2) < G_t^*(t^N(\lambda_1), \lambda_1) = G_t^W(t^N(\lambda_1), \lambda_1)$$

These inequalities and the concavity of G^W in t imply

$$G_t^W(t^{MPA}(\lambda_2), \lambda_2) < G_t^W(t^N(\lambda_2), \lambda_2) < G_t^W(t^N(\lambda_1), \lambda_1) < G_t^W(t^{MPA}(\lambda_1), \lambda_1)$$

This contradicts the FOC, which requires G_t^W to be equalized across states.

Part (ii) can be similarly proved. ■

Proof of Proposition 1:

First observe that $G_{t\lambda} > 0$ implies $t^N(\lambda)$ is increasing, and $G_{t\lambda}^W > 0$ implies $t^{MPA}(\lambda)$ is increasing (this can be proved by implicitly differentiating the FOC for the MPA problem and recalling that ψ is independent of λ).

Part (i). Focus on the case $\frac{d}{d\lambda}G_t^*(t^N(\lambda), \lambda) < 0$. By Lemma 1, in this case $t^{MPA}(\lambda)$ intersects $t^N(\lambda)$ once and from above. We show that the random variable $t^N(\lambda)$ is a second order stochastic shift of the random variable $t^{MPA}(\lambda)$, which together with the fact that these two random variables have the same mean implies that the former is a MPS of the latter. Let $\lambda^N(t)$ denote the inverse of $t^N(\lambda)$ and $\lambda^{MPA}(t)$ the inverse of $t^{MPA}(\lambda)$; these inverse functions exist because $t^N(\lambda)$ and $t^{MPA}(\lambda)$ are both increasing. Also, let $\hat{t}$ be the value of t for which the two curves intersect.

The cdf of t^N is given by $F_N(t) = Pr(t^N(\lambda) \leq t) = Pr(\lambda \leq \lambda^N(t))$ and the cdf of

t^{MPA} is given by $F_{MPA}(t) = Pr(t^{MPA}(\lambda) \leq t) = Pr(\lambda \leq \lambda^{MPA}(t))$. Lemma 1 implies that $\lambda^{MPA}(t) < \lambda^{N}(t)$ for all $t < \hat{t}$ and $\lambda^{MPA}(t) > \lambda^{N}(t)$ for all $t > \hat{t}$, which in turn implies that $F_{MPA}(t) < F_{N}(t)$ for all $t < \hat{t}$ and $F_{MPA}(t) > F_{N}(t)$ for all $t > \hat{t}$. This implies that $t^{N}(\lambda)$ is a second order stochastic shift of $t^{MPA}(\lambda)$, as claimed.

Part (ii) can be similarly proved.

Part (iii) was already proved in Lemma 1. ■

Proof of Proposition 2:

Start by noting that $G^*_{tt} = \frac{\partial^2 G^*}{\partial(\ln p^*)^2}$. It is straightforward to derive:

$$\frac{\partial^2 G^*}{\partial\left(\ln p^*\right)^2} = \left(v^{*\theta}\right)\left[\theta\left(\frac{\partial \ln v^*}{\partial \ln p^*}\right)^2 + \frac{\partial^2 \ln v^*}{\partial\left(\ln p^*\right)^2}\right],$$

where $\ln v^* = \ln R^* - \ln \phi^*$. Next note that $\frac{\partial \ln R^*}{\partial \ln p^*} = \frac{p^* q^*}{R^*}$. Differentiating this elasticity with respect to $\ln p^*$ and simplifying, we obtain:

$$\frac{\partial^2 \ln R^*}{\partial\left(\ln p^*\right)^2} = \frac{p^* q^*}{R^*} \cdot \left(1 - \frac{p^* q^*}{R^*}\right) + \frac{{p^*}^2 {q^*}'}{R^*}.$$

Next note that employing Roy's identity we obtain $\frac{c^*}{R^*} = \frac{{\phi^*}'}{\phi^*}$, hence $\frac{\partial \ln \phi^*}{\partial \ln p^*} = \frac{p^* c^*}{R^*}$. It follows that

$$\frac{\partial^2 \ln \phi^*}{\partial\left(\ln p^*\right)^2} = \frac{\partial\left(\frac{p^* c^*}{R^*}\right)}{\partial p^*} \cdot p^*.$$

Adding things up and simplifying, we find $G^*_{tt} = v^{*\theta}\Omega^*\left(\theta\Omega^* + \varepsilon^*_x + D^*\right)$. ■

Proof of Proposition 4:

We start by proving part (b). The schedules $t^{MPA}(\lambda)$ and $t^{N}(\lambda)$ are clearly continuous. The mean constraint and the continuity of $t^{MPA}(\lambda)$ and $t^{N}(\lambda)$ ensure the existence of at least one intersection. Consider one such intersection $\hat{\lambda}$, so that $t^{MPA}(\hat{\lambda}) = t^{N}(\hat{\lambda})$. By the FOC, $G^W_t(t^N(\hat{\lambda}), \hat{\lambda}, k^{*MPA}) = \psi$. Since $G_t(t^N(\hat{\lambda}), \hat{\lambda}) = 0$ this implies $G^*_t(t^N(\hat{\lambda}), \hat{\lambda}, k^{*MPA}) = \psi$. Now if $\frac{d}{d\lambda}G^*_t(t^N(\lambda), \lambda, k^{*MPA}) = 0$ then $G^*_t(t^N(\lambda), \lambda, k^{*MPA}) = \psi$ for all λ, which in turn implies $G^W_t(t^N(\lambda), \lambda, k^{*MPA}) = \psi$ for all λ. Therefore the schedule $t^N(\lambda)$ satisfies the FOC, hence $t^{MPA}(\lambda) = t^N(\lambda)$ for all λ and $k^{*MPA} = k^{*N}$.

We next prove part (a). Again, $t^{MPA}(\lambda)$ and $t^{N}(\lambda)$ must intersect at least once. We now argue that if $\frac{d}{d\lambda}G^*_t(t^N(\lambda), \lambda, k^{*MPA}) < 0$ for all λ then $t^{MPA}(\lambda)$ can only intersect $t^N(\lambda)$ from above. This, together with continuity, will also ensure the uniqueness of the intersection.

We argue by contradiction. Suppose $t^{MPA}(\lambda)$ intersects $t^N(\lambda)$ at some point $\hat{\lambda}$ from below. Consider two values of λ on the opposite sides of this intersection, $\lambda_1 < \hat{\lambda} < \lambda_2$,such that $t^{MPA}(\lambda_1) < t^N(\lambda_1)$ and $t^{MPA}(\lambda_2) > t^N(\lambda_2)$.

Recalling that $G_t(t^N(\lambda), \lambda) = 0$ for all k^* and assuming $\frac{d}{d\lambda} G_t^*(t^N(\lambda), \lambda, k^{*MPA}) < 0$ for all λ then

$$\begin{aligned} G_t^W(t^N(\lambda_2), \lambda_2, k^{*MPA}) &= G_t^*(t^N(\lambda_2), \lambda_2, k^{*MPA}) \\ &< G_t^*(t^N(\lambda_1), \lambda_1, k^{*MPA}) = G_t^W(t^N(\lambda_1), \lambda_1, k^{*MPA}) \end{aligned}$$

These inequalities and the concavity of G^W in t imply

$$\begin{aligned} G_t^W(t^{MPA}(\lambda_2), \lambda_2, k^{*MPA}) &< G_t^W(t^N(\lambda_2), \lambda_2, k^{*MPA}) \\ &< G_t^W(t^N(\lambda_1), \lambda_1, k^{*MPA}) < G_t^W(t^{MPA}(\lambda_1), \lambda_1, k^{*MPA}). \end{aligned}$$

The claim follows. ■

Proof of Proposition 5

As a first step, we argue that an increase in k^* leads to a decrease in the degree of diversification D^*. We can write $D^* = 1 - \frac{p^*q^*}{p^*q^*+q_0^*} = 1 - \frac{1}{1+\frac{q_0^*}{p^*q^*}}$. An increase in k^* (holding $\pi^* = \ln p^*$ constant) leads to an increase in q^* and a decrease in q_0^*, hence D^* falls.

Next focus on Ω^*. We have $\Omega^* = \frac{p^*x^*}{R^*} = \frac{p^*q^*-p^*c^*}{R^*} = \frac{1}{1+\frac{q_0^*}{p^*q^*}} - \frac{p^*c^*}{R^*}$. As k^* increases, the first term in the above expression increases, as we argued above. Next note that k^* affects the consumption share $\frac{p^*c^*}{R^*}$ only through R^*. In principle $\frac{\partial R^*}{\partial k^*}$ has an ambiguous sign, but note that under certainty k^* maximizes R^*, hence $\frac{\partial R^*}{\partial k^*} = 0$ under certainty. If p^* is uncertain but has a small support, $\frac{\partial R^*}{\partial k^*}$ will be small in absolute value, and hence $\frac{\partial}{\partial k^*}\left(\frac{p^*c^*}{R^*}\right)$ will also be small in absolute value. This ensures that if the support is small enough, Ω^* is increasing in k^*.

Next note that a change in k^* in general has an ambiguous effect on the export supply elasticity ε_x^*, so in general the effect of k^* on $\Omega^*\left(\theta\Omega^* + \varepsilon_x^* + D^*\right)$ is ambiguous, however if risk aversion is sufficiently strong, i.e. if θ is sufficiently negative, then clearly the effect is negative. If ε_x^* is approximately constant we do not require θ to be sufficiently negative.

Finally, consider the sign of the whole expression (VI.4). Letting $\Omega^*\left(\theta\Omega^* + \varepsilon_x^* + D^*\right) \equiv h(p^*, k^*)$, we can rewrite (VI.4) as

$$\text{(X.1)} \qquad \frac{\partial}{\partial k^*}\left[v^{*\theta}(p^*, k^*)h(p^*, k^*)\right] = \frac{\partial v^{*\theta}}{\partial k^*} \cdot h + v^{*\theta} \cdot \frac{\partial h}{\partial k^*} = \left(\theta\frac{v_{k^*}^*}{v^*} + \frac{h_{k^*}}{h}\right) \cdot h \cdot v^{*\theta}$$

Note that the term $\frac{v^*_{k^*}}{v^*}$ is the relative change in real income due to a capital re-allocation. This is zero under certainty, and under uncertainty it necessarily changes sign over the range of k^*, since if it was always positive or negative there would be an incentive to re-allocate capital. We now argue that if θ is sufficiently negative and the support of p^* is small enough, the expression above is negative. Fix θ at some level $\hat{\theta}$ such that $h < 0$ and $\frac{h_{k^*}}{h} > A > 0$ under certainty (where A is some positive constant). The arguments above ensure that such $\hat{\theta}$ must exist. Next recall that k^* satisfies $v^*_{k^*} = 0$ under certainty. Then, as the support of p^* shrinks to zero, $\hat{\theta}\frac{v^*_{k^*}}{v^*}$ goes to zero for all p^* in the support, while $\frac{h_{k^*}}{h}$ approaches $A > 0$, therefore $\frac{\partial}{\partial k^*}\left[v^{*\theta}(p^*, k^*)h(p^*, k^*)\right] < 0$. ■

Proof of Proposition 6:

Focus on the case $(G^{*N}_{tt} + G^{*N}_{tt^*})\frac{dt^N}{d\lambda} + G^{*N}_{t\lambda} < 0$, or equivalently $\frac{d}{d\lambda}G^*_t(t^N(\lambda), t^N(\lambda), \lambda) < 0$. The key is to prove the analog of Lemma 1, namely that $t^{MPA}(\lambda)$ intersects $t^N(\lambda)$ once and from above.

We argue by contradiction. Suppose $t^{MPA}(\lambda)$ intersects $t^N(\lambda)$ at some point $\hat{\lambda}$ from below. Consider two values of λ on the opposite sides of this intersection, $\lambda_1 < \hat{\lambda} < \lambda_2$,such that $t^{MPA}(\lambda_1) > t^N(\lambda_1)$ and $t^{MPA}(\lambda_2) > t^N(\lambda_2)$.

Recalling that $G_t(t^N(\lambda), t^N(\lambda), \lambda) = 0$ and $\frac{d}{d\lambda}G^*_t(t^N(\lambda), t^N(\lambda), \lambda) < 0$ for all λ, then

$$\tilde{G}_t(t^N(\lambda_2), \lambda_2) = G^*_t(t^N(\lambda_2), t^N(\lambda_2), \lambda_2) < G^*_t(t^N(\lambda_1), t^N(\lambda_1), \lambda_1) = \tilde{G}_t(t^N(\lambda_1), \lambda_1)$$

These inequalities and the concavity of $\tilde{G}$ in t imply

$$\tilde{G}_t(t^{MPA}(\lambda_2), \lambda_2) < \tilde{G}_t(t^N(\lambda_2), \lambda_2) < \tilde{G}_t(t^N(\lambda_1), \lambda_1) < \tilde{G}_t(t^{MPA}(\lambda_1), \lambda_1)$$

This contradicts the FOC, which requires that $\tilde{G}_t(t^{MPA}(\lambda), \lambda)$ be equalized across states.

Having proved the analog of Lemma 1, the claim of the proposition follows immediately: just observe that the assumed single crossing properties imply $t^N(\lambda)$ and $t^{MPA}(\lambda)$ are increasing, and apply a similar argument to that in the proof of Proposition 1. ■

Figure 1
Average Tariff Distributions: 1865-1913

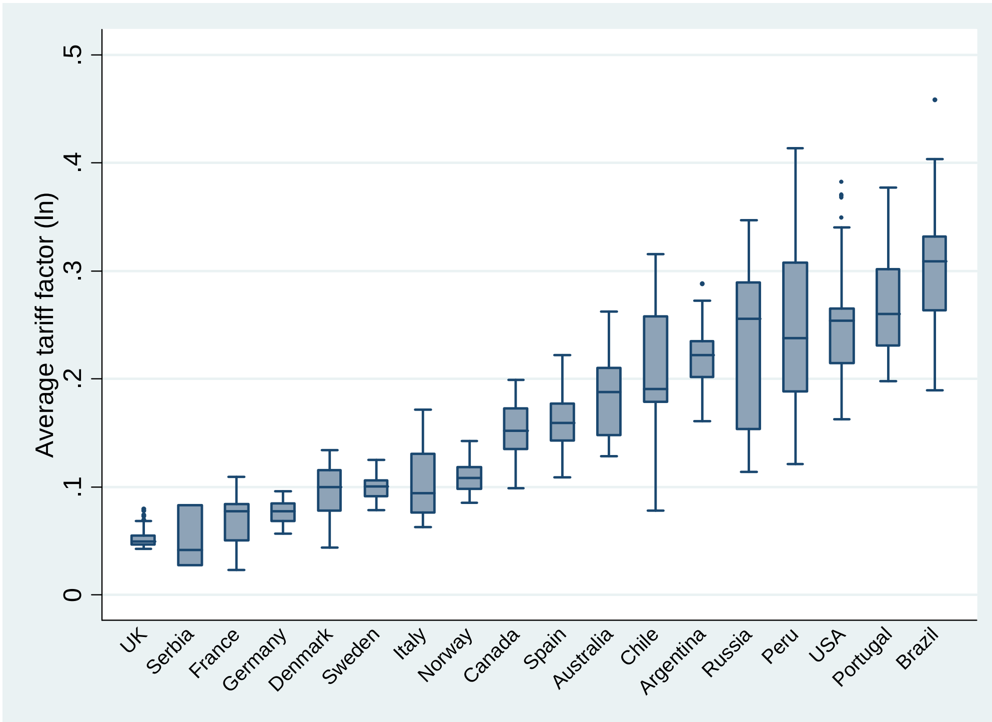

Notes: Data source Schularick and Solomou (2011). Import weighted tariff τ, modified to t=ln(1+τ).

Figure 2

US Average Tariff: 1867-1961

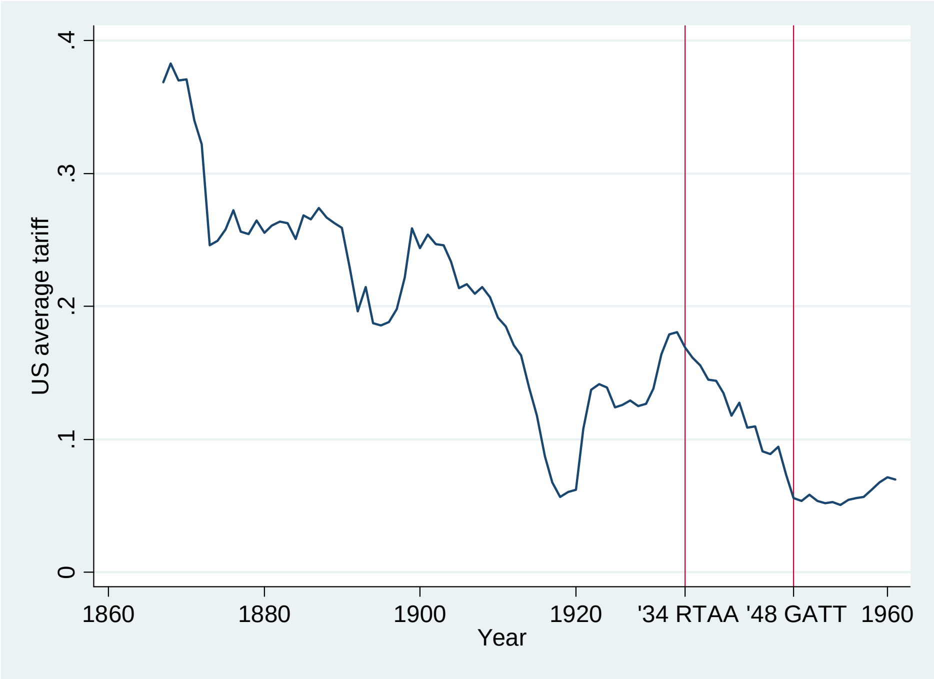

Notes: Data source Irwin (2007) import weighted tariff τ, modified to t=ln(1+τ).

Figure 3
Noncooperative Policy vs. Mean Preserving Agreement

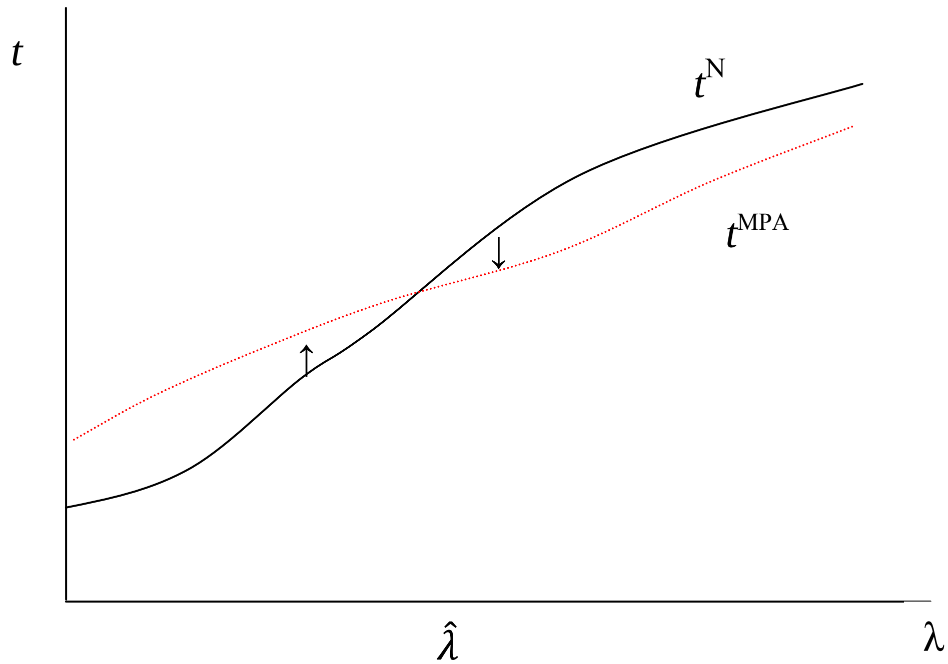

Table 1

Relative gains from reducing tariff uncertainty [1]

		US average tariff (t)	US sugar tariff (t)
	θ		
	0	-0.046	-0.079
	-1	-0.004	0.000
	-2	0.04	0.08
	-3	0.08	0.16
log Cuban adjusted openness	-4	0.12	0.24
($\theta \ln v^* + \ln \Omega^*$)	-5	0.16	0.32
	-6	0.20	0.40
	-7	0.25	0.48
	-8	0.29	0.56
	-9	0.33	0.64
	-10	0.37	0.72

[1] Authors' calculations. Cov(y^*,t)/E(t) for 1903-33 where y* is log(Cuban adjusted openness) and t is either ln(1+tariff) averaged over all products for US or only its tariff on Cuban sugar. See text for data sources. The relative gain is the absolute value of Cov(y^*,t)/E(t).