

Trade Policy Ambiguity

Joseph B. Steinberg*

October 2026

Abstract

This paper studies how Knightian uncertainty about the distribution of future trade policies affects current trade flows using a dynamic trade model with a sunk cost of exporting. Qualitatively, trade-policy ambiguity reduces export participation in a similar manner to standard trade-policy risk, but also increases sensitivity to sunk entry costs and tariff bounds, and can either increase or decrease sensitivity to the likelihood and persistence of tariff increases. Quantitatively, ambiguity dampens the response of trade to persistent reforms and can strengthen the response to transitory reforms.

JEL Classifications: F12, F13, F17

Keywords: Trade-policy uncertainty, tariff risk, ambiguity, Knightian uncertainty, export participation

*joseph.steinberg@utoronto.ca, University of Toronto and NBER.

1 Introduction

Uncertainty about future trade policy affects current trade flows by discouraging export participation. The trade-policy uncertainty literature has focused on historical episodes featuring well-defined tariff schedules and risk over which schedule exporters will face. In the case of U.S.-China trade during the 20th century, for example, the risk was that the United States would revoke China’s Normal Trade Relations status, which would have switched it to the Non-Normal Trade Relations schedule ([Pierce and Schott, 2016](#); [Handley and Limão, 2017](#); [Alessandria et al., 2025b](#)). Starting in 2018, however, the United States has imposed tariff increases on geopolitical foes and close allies alike that depart markedly from the tariff schedules codified in U.S. law, increasing the risk of future tariff changes but also uncertainty about what shape those changes could take. This paper studies how trade-policy ambiguity—Knightian uncertainty over future tariff distributions—affects export participation and aggregate trade.

I begin with empirical evidence consistent with trade-policy ambiguity after the myriad tariff announcements at the start of Donald Trump’s second term in 2025. Surveys show that the range of tariffs firms considered plausible widened sharply after these announcements; tariff forecast dispersion increased both within and across firms. Trade data show that U.S. imports fell more where the gap between the statutory tariff and the worst-case tariff—the rate implied by the most pessimistic possible interpretation of the announcements—was larger, consistent with firms putting more weight on worst-case outcomes in their decision making.

In the theoretical part of the paper, I introduce trade-policy ambiguity into a [Handley and Limão \(2017\)](#) model with sunk entry costs using the robust-control framework of [Hansen and Sargent \(2008\)](#). Firms treat the tariff transition matrix as an approximation and choose decision rules that perform well under perturbations of it.¹ Like an increase in tariff risk, ambiguity reduces export participation, but it also makes participation more sensitive to sunk entry costs and tariff bounds, and can increase or decrease sensitivity to the likelihood and persistence of tariff increases.

In the quantitative part of the paper, I use the dynamic trade model of [Alessandria et al. \(2021\)](#) to study the implications of policy ambiguity for aggregate trade dynamics. I simulate the model with different levels of tariff persistence and ambiguity and estimate trade-elasticity dynamics on each sample using the local projections method of [Boehm et al. \(2023\)](#). Without ambiguity, trade responds more in the long run to persistent tariff changes than transitory ones as shown by [Alessandria et al. \(2025a\)](#). When ambiguity is severe, trade responds similarly regardless of persistence. When it is mild, trade responds less to

¹The appendix studies the multiple-prior framework of [Gilboa and Schmeidler \(1989\)](#), which yields similar effects on entry thresholds, their sensitivity to sunk costs and tariff bounds, and aggregate trade dynamics.

persistent tariff changes but can respond more to transitory ones.

This paper contributes to the trade-policy uncertainty literature by studying ambiguity about the tariff process alongside risk within a known process (Handley, 2014; Handley and Limão, 2017, 2022; Alessandria et al., 2025b,c). It also contributes to the literature on tariff bindings, which cap tariffs while permitting discretion below the bound. Amador and Bagwell (2013) show that such caps are optimal when governments are privately informed about the political pressure they face to protect import-competing industries. Beshkar et al. (2015) document that applied tariffs often lie well below their bounds, leaving governments substantial room to raise them, while Jakubik and Piermartini (2023) find that bindings constrain tariff increases following import shocks. Most closely related, Handley (2014) shows that bindings promote export entry by reducing trade-policy uncertainty. I show that ambiguity makes entry more sensitive to these bounds than standard risk does, so changes in bounds have larger effects on trade.

The paper also contributes to the trade-elasticity measurement literature by showing how ambiguity can help explain variation in estimated elasticities across time horizons and policy environments (Soderbery, 2018; Simonovska and Waugh, 2014; Boehm et al., 2023; Alessandria et al., 2025a). More broadly, this paper contributes to the literature on ambiguity and real-option decision problems. Outside trade, several papers have studied how ambiguity affects irreversible-investment and optimal-stopping problems (Nishimura and Ozaki, 2007; Thijssen, 2011; Huang and Yu, 2021). This is the first paper to study how ambiguity affects export participation.

2 Motivating evidence

In this paper, firms that are uncertain about the distribution of future tariffs put more weight on worst-case outcomes in making decisions about export participation. Two pieces of evidence from 2025 support this interpretation.

First, firms’ views about what tariffs they would face became far less precise as policy became more erratic. The Federal Reserve Bank of Boston surveyed small and medium-sized U.S. businesses about the tariffs they expected the United States to impose in 2025, eliciting expected, minimum, and maximum values from each respondent. Among importers, the average within-respondent standard deviation implied by these forecasts rose from 5.8 percentage points in December 2024 to 10.3 in February 2025 and 13.4 in April 2025. The survey’s authors describe the April 2025 reading as reflecting “heightened trade-policy ambiguity” (Andrade et al., 2025c).

Expectations also varied widely across firms. In December 2024, before any of the new tariffs took effect, roughly one in ten importing firms expected tariffs on Chinese goods below 5%, while a similar share expected them to exceed 50%. Firms that expected higher tariffs

also considered a wider range of tariffs plausible: the gap between minimum and maximum tariff forecasts grew roughly in proportion to expected forecasts (Andrade et al., 2025a). By April 2025, the share of firms expecting tariffs on Chinese goods above 50% had risen from 10% to 30%, but about one in five still expected tariffs below 10% (Andrade et al., 2025b). Canadian firms disagreed even about the probability that tariffs would be imposed: in the Bank of Canada’s Business Outlook Survey for the first quarter of 2025, firms’ assessments of the likelihood that U.S. tariffs would take effect were “distributed across the full range of probability, from 0% to 100%” (Bank of Canada, 2025).

Second, trade appears to have responded to the worst-case tariffs that firms could have inferred from policy announcements. Manova et al. (2026) document what they call “tariff confusion:” U.S. tariff announcements in 2025 were so frequent and complex that, by December, firms could have arrived at 79 different tariff calculations for the average origin country and product. The authors define the worst-case tariff as the tariff a firm would have calculated had it tracked every increase and none of the reductions, and calculate that it was about twice the statutory tariff on average. They show that holding the statutory tariff fixed, country-product pairs with higher worst-case tariffs saw larger declines in imports, and that the elasticity of imports with respect to the gap between the worst-case and statutory tariffs was of the same order of magnitude as the elasticity with respect to the statutory tariff itself. Their focus is confusion about current, not future, tariffs, but the behavior they document is consistent with ambiguity aversion: when firms cannot determine which of several possible tariff schedules applies, their decision-making puts more weight on unfavorable schedules.

3 Theory

I begin by analytically characterizing the effects of trade-policy ambiguity on export participation using the stylized economic model of Handley and Limão (2017). I consider two expectational setups: standard trade-policy risk where the distribution of future tariffs is known with certainty; and trade-policy ambiguity where this distribution is uncertain, modeled using the robust-control framework of Hansen and Sargent (2008).

3.1 Environment

There is a unit measure of firms that are heterogeneous in productivity $z \sim F(z)$, which is exogenous, and export participation, which is endogenous. Firms begin their lives as non-exporters, die with probability $1 - \delta$, and discount the future with a factor β . Exporting firms earn per-period profits

$$\pi_s(z) = Dz\tau_s^{1-\sigma} \equiv zT_s, \quad (1)$$

where D is a constant aggregate demand term, and $\tau_s \geq 1$ is an ad-valorem tariff that depends on the aggregate state s , which I describe below, and σ is the price elasticity of

foreign demand for a firm's product. Non-exporters must pay a one-time sunk cost κ to become exporters. For notational brevity, I define $T_s = D\tau_s^{1-\sigma}$ as the profits of a firm with productivity one.

Tariffs follow a Markov process with two states, $s \in \{L, H\}$ with $\tau_L < \tau_H$, and transition matrix

$$P = \begin{bmatrix} 1 - \eta & \eta \\ 1 - \rho & \rho \end{bmatrix}, \quad (2)$$

where $\eta = \Pr(s' = H|s = L)$ is the hazard into high tariffs and $\rho = \Pr(s' = H|s = H)$ is the persistence of high tariffs. For future reference, the unconditional probability of the high-tariff state is

$$p_H = \frac{\eta}{\eta + 1 - \rho}. \quad (3)$$

3.2 Firm's problem

Non-exporters choose whether to pay the sunk cost and begin exporting. Their problem is described by the Bellman equations:

$$V_s^X(z) = \pi_s(z) + \delta\beta\mathcal{A}_s[V_{s'}^X(z)], \quad (4)$$

$$V_s^N(z) = \max\{\pi_s(z) - \kappa + \delta\beta\mathcal{A}_s[V_{s'}^X(z)], \delta\beta\mathcal{A}_s[V_{s'}^N(z)]\}. \quad (5)$$

$V_s^X(z)$ is the value of starting the period as an exporter in state s , $V_s^N(z)$ is the value of starting as a non-exporter, and $\mathcal{A}$ is the one-step certainty-equivalence operator specific to a given expectational setup.

As in [Handley and Limão \(2017\)](#), entry is characterized by productivity thresholds $z_L^* < z_H^*$ at which the values of entering and waiting are equal:

$$\pi_s(z_s^*) - \kappa + \delta\beta\mathcal{A}_s[V_{s'}^X(z_s^*)] = \delta\beta\mathcal{A}_s[V_{s'}^N(z_s^*)]. \quad (6)$$

Firms with productivity below z_L^* never enter, those between the thresholds enter only when tariffs are low, and those above z_H^* enter in either state. Thus, $V_s^N(z) = V_s^X(z) - \kappa$ in both states for $z \geq z_H^*$, which implies that the high-state entry condition can be written as

$$\pi_H(z_H^*) - \kappa + \delta\beta\mathcal{A}_H[V_{s'}^X(z_H^*)] = \delta\beta\mathcal{A}_H[V_{s'}^X(z_H^*) - \kappa]. \quad (7)$$

In both setups considered below, $\mathcal{A}$ is translation invariant: $\mathcal{A}[V + c\mathbf{1}] = \mathcal{A}[V] + c$. The continuation values therefore cancel from the high-state condition, yielding the following

solution for the high-tariff threshold regardless of the expectational setup:²

$$z_H^* = \frac{\kappa(1 - \delta\beta)}{T_H}. \quad (8)$$

The condition that characterizes the low-tariff entry threshold can also be simplified by using the fact that the value of starting as a non-exporter is zero below this threshold: $V_s^N(z) = 0$ for $s \in \{L, H\}$ and $z \leq z_L^*$. This implies that at z_L^* , the sunk cost is exactly equal to the value of starting as an exporter:

$$\kappa = V_L^X(z_L^*) = \pi_L(z_L^*) + \delta\beta\mathcal{A}_L[V_{s'}^X(z_L^*)]. \quad (9)$$

The remaining analysis studies how ambiguity changes this condition.

3.3 Solution with standard policy risk

The standard assumption in the trade-policy uncertainty literature is that while firms may not have perfect foresight over tariffs, they do know the true transition matrix P .³ I refer to this setup as trade-policy risk. In this case, the certainty-equivalence operator is simply the expected value:

$$\mathcal{A}_s^{SR}[V_{s'}^X] = \mathbb{E}_P[V_{s'}^X] = \sum_{s'=L,H} P_{s,s'} V_{s'}^X. \quad (10)$$

To solve for z_L^* , we need to know the value of $V_L^X(z_L^*) = \pi_L(z_L^*) + \delta\beta\mathcal{A}_L[V_{s'}^X(z_L^*)]$. Matrix form is the easiest way to write this:

$$\begin{bmatrix} V_L^X(z_L^*) \\ V_H^X(z_L^*) \end{bmatrix} = \begin{bmatrix} \pi_L(z_L^*) \\ \pi_H(z_L^*) \end{bmatrix} + \delta\beta P \begin{bmatrix} V_L^X(z_L^*) \\ V_H^X(z_L^*) \end{bmatrix} = (I - \delta\beta P)^{-1} \begin{bmatrix} z_L^* T_L \\ z_L^* T_H \end{bmatrix}. \quad (11)$$

The value for the low state can be expressed in a more economically intuitive fashion as

$$V_L^X(z_L^*) = z_L^* \left[\frac{\bar{T}}{1 - \beta\delta} + \frac{T_L - \bar{T}}{1 - \beta\delta(\rho - \eta)} \right], \quad (12)$$

where $\bar{T} = p_H T_H + (1 - p_H) T_L$. Thus, (9) simplifies to

$$z_L^{*,SR} = \frac{\kappa}{\left[\frac{\bar{T}}{1 - \beta\delta} + \frac{T_L - \bar{T}}{1 - \beta\delta(\rho - \eta)} \right]}. \quad (13)$$

²This result need not hold with continuation costs and stochastic productivity, which the quantitative model introduces.

³Studies that allow the policy transition matrix to vary over time, such as [Alessandria et al. \(2025b,c\)](#), either assume that firms know the entire path of transition matrices with certainty, or that they know the current matrix at any given point in time and believe it will last forever.

The properties of the threshold z_L^* under standard trade-policy risk are intuitive and well-known. It is increasing in the sunk entry cost, κ , which makes entry less attractive, and it is increasing in the high-tariff value, τ_H , the probability of switching from the low-tariff state to the high-tariff state, η , and the probability of staying in the high-tariff state, ρ , all of which increase the expected present value of tariffs when starting in the low-tariff state.

3.4 Solution with policy ambiguity

Under trade-policy ambiguity, firms are uncertain not only about the path of future tariffs, but also about the transition matrix P that governs the stochastic process that tariffs follow. I use the robust-control framework of [Hansen and Sargent \(2008\)](#), in which firms treat P as an approximation and form decision rules that perform well under perturbations of P . This idea can be formalized by positing a two-player problem in which “nature” seeks to minimize the firm’s payoff by perturbing the transition matrix but pays a penalty proportional to the relative entropy between its choice and the true matrix.⁴

Decision theorists distinguish between ambiguity about which of several candidate models is correct and concern that all candidates are misspecified ([Hansen and Marinacci, 2016](#)). The robust-control framework captures the latter. In the appendix, I analyze the multiple-prior approach of [Gilboa and Schmeidler \(1989\)](#), which captures the former, and find similar main results. I use the term “trade-policy ambiguity” broadly to mean Knightian uncertainty about the stochastic process for tariffs, as opposed to risk within a known process. This is consistent with the macroeconomics literature, which typically treats robust control as one of several representations of ambiguity aversion (see, e.g., [Huo et al., 2024](#)), and with decision theory, in which both robust-control and multiple-prior preferences are special cases of the variational preferences of [Machcheroni et al. \(2006\)](#) ([Strzalecki, 2011](#)).

The one-step operator under robust control is

$$\mathcal{A}_s^{RC} [V_{s'}^X; \theta] = -\theta \log \left[\sum_{s'=L,H} P_{s,s'} e^{-V_{s'}^X/\theta} \right]. \quad (14)$$

The parameter θ is the entropy penalty. When θ is sufficiently large, this operator can be approximated by a second-order expansion:

$$\mathcal{A}_s^{RC} [V_{s'}^X; \theta] \approx \mathbb{E}_P [V_{s'}^X | s] - \frac{1}{2\theta} \text{Var}_P [V_{s'}^X | s]. \quad (15)$$

⁴The same functional form as (14) arises in the “smooth ambiguity” model of [Klibanoff et al. \(2005\)](#) when ambiguity aversion takes an exponential form. There, however, it represents robustness across priors over a set of models, rather than robustness to misspecification of a single model ([Hansen and Marinacci, 2016](#)). [Huo et al. \(2024\)](#) show that, despite this conceptual difference, the two approaches can yield equivalent equilibrium strategies in a Gaussian environment, but this equivalence does not hold in general.

The value of starting as an exporter in the low-tariff state can then be approximated as a quadratic function of z_L^* ,

$$V_L^X(z_L^*) \approx b_L z_L^* - \frac{a_L}{2\theta} (z_L^*)^2, \quad (16)$$

where

$$\begin{bmatrix} b_L \\ b_H \end{bmatrix} := (I - \delta\beta P)^{-1} \begin{bmatrix} T_L \\ T_H \end{bmatrix}, \quad \begin{bmatrix} a_L \\ a_H \end{bmatrix} := \delta\beta(I - \delta\beta P)^{-1} \begin{bmatrix} \eta(1-\eta)(b_H - b_L)^2 \\ \rho(1-\rho)(b_H - b_L)^2 \end{bmatrix}. \quad (17)$$

The coefficient b_L is the value of exporting per unit of productivity under standard risk, so $z_L^{*,SR} = \kappa/b_L$. The low-state threshold is the solution to

$$\frac{a_L}{2\theta} (z_L^*)^2 - b_L z_L^* + \kappa = 0, \quad (18)$$

which is given by

$$z_L^{*,RC}(\theta) = \frac{b_L - \sqrt{b_L^2 - 2a_L\kappa/\theta}}{a_L/\theta}. \quad (19)$$

This root exists for $\theta \geq 2a_L\kappa/b_L^2$, in which case $z_L^{*,RC} \leq 2z_L^{*,SR}$; the approximation applies only when ambiguity is not so strong that it more than doubles the entry threshold.

The first-order approximation of this solution is more economically transparent:

$$z_L^{*,RC} \approx \frac{\kappa}{b_L} + \frac{a_L\kappa^2}{2\theta b_L^3}. \quad (20)$$

It shows that the solution under ambiguity coincides with the standard-risk solution when there is no ambiguity (i.e., the entropy penalty θ is infinite), and increases with the degree of ambiguity (i.e., increases when θ falls). However, there is an additional term which is increasing in the degree of ambiguity. Formally,

$$z_L^{*,RC} > z_L^{*,SR}, \quad \frac{\partial z_L^{*,RC}}{\partial \theta} < 0. \quad (21)$$

3.5 Comparative statics

Ambiguity reduces export participation, but also changes how export participation responds to entry costs and the tariff process. The proofs of the results below, which are relegated to the appendix, use the quadratic approximation (16) unless noted otherwise. The appendix also solves the exact model over a wide range of parameter values, including low values of θ (i.e., high degrees of ambiguity) where the approximation does not apply. All results below hold in every case except the first-order sign conditions for η and ρ , whose accuracy declines as ambiguity becomes stronger.

3.5.1 Sunk entry cost (κ)

Higher sunk costs raise the entry threshold in both setups:

$$\frac{\partial z_L^{*,SR}}{\partial \kappa} > 0, \quad \frac{\partial z_L^{*,RC}}{\partial \kappa} > 0. \quad (22)$$

Because policy ambiguity reduces the certainty-equivalent of export profits, it makes the entry threshold more sensitive to the sunk cost:

$$\frac{\partial z_L^{*,RC}}{\partial \kappa} > \frac{\partial z_L^{*,SR}}{\partial \kappa}. \quad (23)$$

A corollary is that a marginal increase in the sunk cost has a larger impact on exporting when ambiguity is stronger:

$$\frac{\partial^2 z_L^{*,RC}}{\partial \theta \partial \kappa} < 0. \quad (24)$$

Under standard risk, the threshold increases linearly with the sunk cost. However, there is an additional quadratic term under policy ambiguity:

$$\frac{\partial^2 z_L^{*,SR}}{\partial \kappa^2} = 0, \quad \frac{\partial^2 z_L^{*,RC}}{\partial \kappa^2} > 0. \quad (25)$$

This is because robust control penalizes the variance in future payoffs, whereas the firm's objective in the standard risk setup is linear in the payoff stream. As the marginal entrant gets more productive—which is what happens when κ rises—the gap between the marginal entrant's payoffs under low and high tariffs rises, and this amplifies the cost of ambiguity under robust control.⁵ Thus, like trade-policy risk in [Handley and Limão \(2017\)](#), ambiguity depresses trade more in industries with larger sunk costs. Section 3.6 shows that, unlike risk, it does so proportionally more.

3.5.2 High tariff value (τ_H)

A higher tariff bound raises the entry threshold in both setups:

$$\frac{\partial z_L^{*,SR}}{\partial \tau_H} > 0, \quad \frac{\partial z_L^{*,RC}}{\partial \tau_H} > 0. \quad (26)$$

Trade-policy ambiguity makes the potential increase in tariffs more important:

$$\frac{\partial z_L^{*,RC}}{\partial \tau_H} > \frac{\partial z_L^{*,SR}}{\partial \tau_H}. \quad (27)$$

⁵The appendix proves (24)–(25) in the exact model for every finite $\theta > 0$. The threshold becomes linear in κ only in the limit as θ goes to zero.

This is because a greater potential tariff increase reduces the mean profits from exporting, $\mathbb{E}_P [V_{s'}^X | s]$, but also increases the variance penalty, $\frac{1}{2\theta} \text{Var}_P [V_{s'}^X | s]$. As with the sunk cost, changes in the value of the high tariff have more impact on exporting when the degree of ambiguity is greater:

$$\frac{\partial^2 z_L^{*,RC}}{\partial \theta \partial \tau_H} < 0. \quad (28)$$

These results imply that lowering tariff bounds promotes entry more when trade policy is ambiguous. They complement evidence that bindings promote export entry under standard trade-policy risk (Handley, 2014), and suggest that WTO bindings and legislative limits on executive tariff powers matter more when ambiguity is greater. In the United States, for example, such limits could restrict the President's ability to raise tariffs under the Trade Act of 1974, the Trade Expansion Act of 1962, or the International Emergency Economic Powers Act.

3.5.3 Likelihood and persistence of high tariffs (η, ρ)

In both setups, the more likely tariffs are to rise and the more persistent such an increase is, the more productive a firm must be to enter when tariffs are low:

$$\frac{\partial z_L^{*,SR}}{\partial p} > 0, \quad \frac{\partial z_L^{*,RC}}{\partial p} > 0, \quad p = \eta, \rho. \quad (29)$$

These effects hold in the exact, unapproximated model for every $\theta > 0$.⁶ Unlike the effects of the sunk cost and the tariff bound, which ambiguity always amplifies, these effects can be either amplified or attenuated by ambiguity.

Using the first-order approximation (20), one can show that

$$\text{sign} \left(\frac{\partial z_L^{*,RC}}{\partial \eta} - \frac{\partial z_L^{*,SR}}{\partial \eta} \right) = \text{sign} \left(\frac{1}{\eta} - \frac{1 - \gamma\rho}{A} - \frac{3\gamma q}{B} \right), \quad (30)$$

$$\text{sign} \left(\frac{\partial z_L^{*,RC}}{\partial \rho} - \frac{\partial z_L^{*,SR}}{\partial \rho} \right) = \text{sign} \left(\frac{3}{B} - \frac{2\rho - \eta}{A} \right), \quad (31)$$

where $\gamma \equiv \beta\delta$, $q \equiv T_H/T_L \in (0, 1)$, $A \equiv (1 - \gamma\rho)(1 - \eta) + \gamma\rho(1 - \rho)$, and $B \equiv 1 - \gamma\rho + \gamma\eta q$. The degree of ambiguity does not play a role in these conditions, nor does the sunk cost; only the transition probabilities, the discount factor γ , and the ratio of high-tariff to low-tariff profits q do. One can further show that if the right-hand side of (30) is positive, so is the right-hand side of (31). Thus, the (η, ρ) plane is divided into three regions: one where z_L^* is more sensitive to both parameters under policy ambiguity than under standard risk;

⁶Under the quadratic approximation, these signs require an additional restriction on θ . The appendix characterizes this restriction and proves the result in the exact model.

another where z_L^* is more sensitive only to ρ ; and a third where z_L^* is less sensitive to both parameters.

All three regions are nonempty for any values of γ and q . The tariff bound determines where the boundaries between the regions lie: both right-hand sides are decreasing in q , so a higher tariff bound τ_H , which lowers q , expands the region where ambiguity makes z_L^* more sensitive to both parameters and contracts the region where it makes z_L^* less sensitive to both. Figure 1 illustrates this for a low and a high tariff bound. Broadly, ambiguity amplifies the effects of risk when tariff hikes are unlikely and attenuates them when hikes are likely and persistent.

Since these results were derived from the model's first-order approximation, they describe how a small degree of ambiguity alters the effects of trade-policy risk. The appendix shows that the three-region structure survives in the exact model, but that stronger ambiguity shrinks the region where z_L^* is more sensitive to both parameters and expands the region where it is less sensitive to both. This mirrors the max-min setup, also analyzed in the appendix, where the entry threshold becomes insensitive to η and ρ as ambiguity becomes severe.

3.6 Distinguishing ambiguity from risk

Ambiguity and risk both raise the entry threshold and reduce export participation, which implies that a decline in export participation cannot distinguish a situation where firms have become more pessimistic about the tariff process from one where a tariff increase has simply become more likely. The two mechanisms do, however, have different implications for how export participation responds to changes in other parameters.⁷

The first difference concerns the sunk cost. Consider an economy where the sunk entry cost, κ , varies across industries as in [Handley and Limão \(2017\)](#). Under standard risk, the entry threshold is proportional to the sunk cost, $z_L^{*,SR} = \kappa/b_L$, so a change in the tariff process (an increase in either the probability or persistence of a tariff hike) raises the thresholds of all industries by the same proportion.⁸ Under ambiguity, the elasticity of the threshold with respect to the sunk cost can be written as

$$\frac{\partial \log z_L^{*,RC}}{\partial \log \kappa} = 1 - \frac{\partial \log z_L^{*,RC}}{\partial \log \theta} > 1. \quad (32)$$

Thus, relative to standard risk, ambiguity raises entry thresholds proportionally more in industries with larger sunk costs.

⁷I establish these differences in the exact robust-control model described in the appendix, which does not rely on the quadratic or first-order approximations used above.

⁸In absolute terms, an increase in risk also raises the threshold more when the sunk cost is larger, as in [Handley and Limão \(2017\)](#). The distinction here is in proportional terms.

The second difference concerns the tariff bound. Consider the following thought experiment. Begin with a baseline economy with standard risk, and compare it to two counterfactual economies where the entry threshold increases by the same amount, but for different reasons. In the first counterfactual, the threshold increases because a tariff increase becomes more likely or more persistent (it does not matter which, as long as the effect on the threshold is the same). In the second, the threshold increases because firms now face ambiguity. Since the threshold rises by the same amount in both, this observation alone cannot tell them apart. The question, then, is whether trade has become more or less sensitive to subsequent changes in the tariff bound in one counterfactual relative to the other. As I prove in the appendix, the elasticity of the threshold with respect to the bound is higher in the ambiguity counterfactual than the increased-risk counterfactual. The intuition is that firms put extra weight on the high-tariff state under ambiguity and this weight rises with the high tariff, so raising the bound lowers the value of exporting twice: through lower profits in the high-tariff state (which is also true under standard risk), and through greater pessimism about reaching that state.

Table 1 illustrates this idea numerically. In the baseline economy, I set $\eta = 0.1$, $\rho = 0.8$, and $\tau_H = 1.25$. Raising the threshold by 10% through a higher likelihood of a tariff increase ($\eta = 0.224$) or a higher persistence ($\rho = 0.998$) raises the elasticity of the threshold with respect to the bound from 0.38 to 0.73 in both cases. Raising the threshold by the same amount through ambiguity raises this elasticity to 1.15—about 60% more.

These results yield testable implications that distinguish ambiguity from risk. First, an increase in risk should reduce export participation by the same proportion in all industries, whereas ambiguity should reduce it proportionally more in industries with larger sunk costs.⁹ Second, given the same observed drop in export participation, entry should respond more to subsequent changes in the tariff bound if the drop was caused by ambiguity rather than by higher risk.

4 Quantification

I now use a quantitative dynamic trade model to explore the implications of trade-policy ambiguity for trade adjustment dynamics.

⁹Technically, the “same proportion” result is only guaranteed to be true when productivity follows a Pareto distribution. In that case, the share of firms above the threshold is proportional to $(z_L^*)^{-k}$, where k is the shape parameter, so equal proportional increases in thresholds imply equal proportional declines in participation. Under other distributions, the elasticity of participation with respect to the threshold depends on the threshold’s level. Regardless, ambiguity makes participation more sensitive (if not exactly in a proportional sense) to the sunk cost under any productivity distribution.

4.1 Model

The model is a partial-equilibrium version of [Alessandria et al. \(2021\)](#). The key differences relative to the simple model used in the previous section are fixed costs to continue exporting and idiosyncratic shocks to productivity and variable trade costs. Together, these additional ingredients generate endogenous exit and gradual trade growth at both the micro and macro levels.

4.1.1 Production and demand

Firms use labor to produce output according to a constant-returns-to-scale technology:

$$y_t = z_t \ell_t. \quad (33)$$

Productivity evolves over time according to an AR(1) process in logs:

$$\log z_t = \rho_z \log z_{t-1} + \sigma_z \varepsilon_t, \quad (34)$$

where ε_t is i.i.d. across firms and time. Firms produce differentiated goods and compete monopolistically. Foreign demand for a firm's good is a downward-sloping function of the price, p_t , and the import tariff, τ_t :

$$d_t(p_t, \tau_t) = (p_t \tau_t)^{-\sigma}, \quad (35)$$

where σ is the price elasticity of demand.

4.1.2 Trade costs and life cycles

There are three types of trade costs: import tariffs, variable trade costs, and fixed trade costs. The import tariff is an aggregate state variable that follows a symmetric version of the two-state Markov process described in [Section 3.1](#), with $\rho = 1 - \eta$. The appendix shows results for discretized AR(1) processes with more than two states.

Variable trade costs are firm-specific and can take three values, $\infty > \xi_H > \xi_L$. The fixed trade cost is a function of the iceberg cost. To begin exporting, a non-exporter with $\xi = \infty$ must pay a fixed cost $\kappa(\infty) = \kappa_0$. To continue exporting, a low- or high-cost exporter must pay a fixed cost $\kappa(\xi_H) = \kappa(\xi_L) = \kappa_1$.

Firms are born as non-exporters. When a firm begins exporting, its variable trade cost falls to ξ_H in the next period. When a high-cost exporter chooses to continue exporting, its variable cost has a chance $1 - \rho_\xi$ of falling to ξ_L in the next period. Symmetrically, a low-cost exporter retains its value of ξ with probability ρ_ξ . In each period, firms die with probability $1 - \delta(z) = \max[0, \min(e^{-\delta_0 z} + \delta_1, 1)]$, in which case they are replaced by non-exporters with productivities drawn from the ergodic distribution.

4.1.3 Firm's problem

The firm's state variables are its productivity, z , its variable trade cost, ξ , and the tariff, τ . The firm's problem has a static component and a dynamic component. The static problem entails choosing a price to maximize the flow profits from exporting given the current state:

$$\pi(z, \xi, \tau) = \max_p \left\{ pd(p, \tau) - w \frac{\xi d(p, \tau)}{z} \right\}. \quad (36)$$

The dynamic problem entails deciding whether to export in the next period. The value of a firm that chooses to export in $t + 1$ is

$$V^X(z, \xi, \tau) = -\kappa(\xi) + \frac{\delta(z)}{1+r} \mathcal{A}_\tau \{ \mathbb{E}_{z', \xi'} [V(z', \xi', \tau')] \}, \quad (37)$$

and the value of a firm that chooses not to export is

$$V^N(z, \xi, \tau) = \frac{\delta(z)}{1+r} \mathcal{A}_\tau \{ \mathbb{E}_{z'} [V(z', \infty, \tau')] \}, \quad (38)$$

where

$$V(z, \xi, \tau) = \pi(z, \xi, \tau) + \max\{V^N(z, \xi, \tau), V^X(z, \xi, \tau)\}. \quad (39)$$

Note that the firm takes expectations about its idiosyncratic state variables, z and ξ , before applying the one-step operator $\mathcal{A}$; there is no ambiguity about idiosyncratic shocks, only the aggregate tariff state. I solve the model numerically using the exact operator (14), without the approximations used in Section 3. Similar to the simple model in the previous section, the solution to the dynamic problem is characterized by a threshold productivity, $z^*(\xi, \tau)$, but it now depends on the firm's variable trade cost and the tariff state.

4.1.4 Aggregation

Let $\varphi_t(z, \xi)$ denote the distribution of firms over productivity and variable trade costs. This distribution reflects past tariffs through firms' entry and continuation decisions, together with idiosyncratic shocks and firm turnover. The law of motion for the distribution is shown in the appendix. Aggregate exports depend on this distribution and the current tariff:

$$X_t = \sum_{\xi \in \{\xi_L, \xi_H\}} \int_z p(z, \xi, \tau_t) d(p(z, \xi, \tau_t), \tau_t) d\varphi_t(z, \xi). \quad (40)$$

4.1.5 Parameter values

I use the parameter values that [Alessandria et al. \(2025a\)](#) calibrated to match Vietnamese firm-level data. Table 2 below lists these values.

One parameter to take particular note of is σ , the demand elasticity, which is the elasticity of trade to tariffs in the short run. The value for this parameter was originally taken from [Soderbery \(2018\)](#). There is a great deal of disagreement in the empirical literature about the trade elasticity, with recent estimates ranging from less than two ([Boehm et al., 2023](#)) to greater than ten ([Alessandria et al., 2025a](#)). As [Alessandria et al. \(2025a\)](#) show, one potential explanation for this disagreement is that trade responds less to transitory tariff shocks than permanent ones, and while quantitative analyses focus on the latter, trade data consist mostly of the former. In this paper, I build on this idea and show that ambiguity about the tariff process can also help explain the dispersion in empirical estimates.

4.2 Experiments

To study how trade-policy ambiguity affects aggregate trade, I estimate trade-elasticity dynamics using simulated data generated by the model. Following [Alessandria et al. \(2025a,b,c\)](#), I assume now that there is a large number of products indexed by $g = 1, \dots, G$, each with its own continuum of firms operating as described in Section 4.1 above, and that tariff shocks are independent across products but common across all firms producing the same product. For a given parameterization, I simulate the model for a large number of periods $t = 1, \dots, T$, resulting in a panel dataset at the product-time level with $G \times T$ observations.

I estimate the trade elasticities using local projections following [Boehm et al. \(2023\)](#). The empirical specification is

$$\Delta_h X_{gt} = -\beta_h^X \Delta_h \tau_{gt} + \lambda_h^X \Delta_0 X_{g,t-1} + \mu_h^X \Delta_0 \tau_{g,t-1} + \delta_g + \delta_t + u_{gt}, \quad (41)$$

where Δ_h is an operator that takes the log difference between a variable's value in period $t+h$ and its value in period $t-1$, $\Delta_0 X_{g,t-1}$ and $\Delta_0 \tau_{g,t-1}$ are the previous period's log changes in exports and tariffs, and δ_g and δ_t are product and time fixed effects. The h -period tariff change $\Delta_h \tau_{gt}$ is instrumented using the one-period change:

$$\Delta_h \tau_{gt} = \beta_h^\tau \Delta_0 \tau_{gt} + \lambda_h^\tau \Delta_0 X_{g,t-1} + \mu_h^\tau \Delta_0 \tau_{g,t-1} + \delta_g + \delta_t + u_{gt}. \quad (42)$$

I simulate the model and run the estimation procedure separately for different combinations of tariff persistence, ρ , and the degree of ambiguity, θ , to study how β_h^X varies as these parameters change. I set the maximum horizon to 14. Beyond this value, estimates lose stability for extremely persistent reforms because tariffs change very infrequently.

4.3 Results

Figure 2 shows the results. Under standard risk, the long-run response of trade rises steeply with tariff persistence: firms respond little to transitory reforms because they know these

reforms are likely to be reversed (Alessandria et al., 2025a), but respond strongly to persistent ones. When tariffs are persistent ($\rho = 0.95$ and $\rho = 0.99$), ambiguity attenuates the long-run response, because firms act as if tariff reductions are more likely to be reversed than they really are. This attenuation is small when ambiguity is mild (i.e., θ is higher), and when $\rho = 0.95$ trade even responds slightly more than under standard risk in the first several years. When tariffs are transitory ($\rho = 0.8$), the trade response under mild ambiguity is substantially larger than under standard risk at all horizons, but when ambiguity is severe enough the response is smaller.

Figure 2(d) summarizes these results by plotting the ratio of the long-run trade elasticity estimate, β_{14}^X , to the short-run estimate, β_0^X . For less persistent reforms ($\rho = 0.8$), mild trade-policy ambiguity amplifies the long-run trade response, but strong ambiguity attenuates it. For more persistent reforms ($\rho = 0.95$ and $\rho = 0.99$), on the other hand, the long-run trade response is always smaller under trade-policy ambiguity than standard risk. In general, trade-policy ambiguity makes aggregate trade less sensitive to the true level of tariff persistence: the differences in trade responses between low- and high-persistence tariff processes are smaller under trade-policy ambiguity than under standard risk, for both mild and strong degrees of ambiguity. Under mild ambiguity, this compression works from both ends: the response to low-persistence reforms is larger and the response to high-persistence reforms is smaller.

How do these results compare with empirical estimates? Because the short-run elasticity in the model equals σ by construction, regardless of tariff persistence or the degree of ambiguity, I compare ratios of long- to short-run elasticities rather than levels. Empirical estimates of this LR-SR ratio range from about 2.3 to 3.6 (Khan and Khederlarian, 2021; Boehm et al., 2023; Alessandria et al., 2025b). I focus on Boehm et al. (2023), whose estimates are at the lower end of this range, because they report not only trade dynamics, but the dynamics of tariffs themselves. They find that about 75% of a change in MFN tariffs remains in place after ten years, which corresponds to a value of $\rho = 0.966$ in the model.

At this level of persistence, the model’s LR-SR ratio is 3.3 under standard risk, which is about 25% higher than the ratio of about 2.6 implied by Boehm et al. (2023)’s point estimates. How much ambiguity would it take to get the model to match that lower value? I find that the answer is a value of about $\theta = 0.2$; only a moderate degree of ambiguity is required to align the model with their findings. Note that in the model, the LR-SR ratio for a perfectly credible, permanent tariff change is much higher, at about 4.3, in line with Alessandria et al. (2025a). Thus, this level of ambiguity is sufficient to reconcile Alessandria et al. (2025a)’s findings of large trade responses for such hypothetical “canonical” reforms with the smaller responses documented by Boehm et al. (2023) for typical tariff changes—the same model is capable of generating both, provided one allows for a bit of ambiguity.

5 Conclusion

This paper compares the effects of trade-policy risk, where future tariffs are uncertain but their distribution is known, with trade-policy ambiguity, where the distribution itself is uncertain. Theoretically, I show that policy risk and ambiguity both reduce trade, but the latter also makes trade more sensitive to entry costs and tariff bounds, and usually makes trade less sensitive to tariff persistence. Quantitatively, I show that these effects manifest in weaker responses of aggregate trade to persistent tariff changes but potentially stronger responses to transitory changes.

My findings add to the conversation about the magnitude of the trade elasticity by illustrating how ambiguity can reduce measured trade responses and also make trade respond similarly to both transitory and persistent reforms. They also underscore the importance of credible and enforceable rules that limit the size of tariff increases, like WTO bindings ([Handley, 2014](#); [Jakubik and Piermartini, 2023](#)) and legislative constraints on executive tariff powers. These rules and institutions are more important when ambiguity is salient, so strengthening them could have particularly pronounced effects on trade in the current geopolitical environment.

6 Declaration of generative AI use

During the preparation of this work the author used generative AI tools provided by OpenAI and Anthropic to verify mathematical proofs and debug computer code. The author reviewed and edited all generated content and takes full responsibility for the published article.

References

- Alessandria, George, Horag Choi, and Kim J. Ruhl, “Trade adjustment dynamics and the welfare gains from trade,” *Journal of International Economics*, 2021, *131*, 103458.
- , Shafaat Yar Khan, Armen Khederlarian, Kim J. Ruhl, and Joseph B. Steinberg, “Recovering Credible Trade Elasticities from Incredible Trade Reforms,” Working Paper 33568, National Bureau of Economic Research 2025.
- , —, —, —, —, and —, “Trade Policy Dynamics: Evidence from 60 Years of US-China Trade,” *Journal of Political Economy*, 2025, *133* (3), 713–749.
- , —, —, —, —, and —, “Trade war and peace: U.S.-China trade and tariff risk from 2015–2050,” *Journal of International Economics*, 2025, *155*, 104066.
- Amador, Manuel and Kyle Bagwell, “The Theory of Optimal Delegation With an Application to Tariff Caps,” *Econometrica*, 2013, *81* (4), 1541–1599.
- Andrade, Philippe, Alexander M. Dietrich, John Leer, Raphael S. Schoenle, Jenny Tang, and Egon Zakrajšek, “Small and Medium-sized Businesses’ Expectations Concerning Tariffs, Costs, and Prices,” Current Policy Perspectives 25-7, Federal Reserve Bank of Boston 2025.
- , —, —, —, Xiao Lin, Raphael S. Schoenle, Jenny Tang, and Egon Zakrajšek, “Who Will Pay for Tariffs? Businesses’ Expectations about Costs and Prices,” Current Policy Perspectives 25-13, Federal Reserve Bank of Boston 2025.
- , —, —, Sophie Handley, John Leer, Raphael S. Schoenle, Jenny Tang, and Egon Zakrajšek, “Effects of Tariff Uncertainty on the Outlook of Small and Medium-sized Businesses,” Current Policy Perspectives 25-12, Federal Reserve Bank of Boston 2025.
- Bank of Canada, “Business Outlook Survey—First Quarter of 2025,” April 2025.
- Beshkar, Mostafa, Eric W. Bond, and Youngwoo Rho, “Tariff binding and overhang: Theory and evidence,” *Journal of International Economics*, 2015, *97* (1), 1–13.
- Boehm, Christoph E., Andrei A. Levchenko, and Nitya Pandalai-Nayar, “The Long and Short (Run) of Trade Elasticities,” *American Economic Review*, 2023, *113* (4), 861–905.
- Gilboa, Itzhak and David Schmeidler, “Maxmin expected utility with non-unique prior,” *Journal of Mathematical Economics*, 1989, *18* (2), 141–153.

- Handley, Kyle**, “Exporting under trade policy uncertainty: Theory and evidence,” *Journal of International Economics*, 2014, *94* (1), 50–66.
- **and Nuno Limão**, “Policy Uncertainty, Trade, and Welfare: Theory and Evidence for China and the United States,” *American Economic Review*, 2017, *107* (9), 2731–2783.
- **and –**, “Trade Policy Uncertainty,” *Annual Review of Economics*, August 2022, *14* (1), 363–395.
- Hansen, Lars Peter and Massimo Marinacci**, “Ambiguity Aversion and Model Misspecification: An Economic Perspective,” *Statistical Science*, 2016, *31* (4), 511–515.
- **and Thomas J. Sargent**, *Robustness*, Princeton, NJ: Princeton University Press, 2008.
- Huang, Yu-Jui and Xiang Yu**, “Optimal stopping under model ambiguity: A time-consistent equilibrium approach,” *Mathematical Finance*, 2021, *31* (3), 979–1012.
- Huo, Zhen, Marcelo Pedroni, and Guangyu Pei**, “Bias and Sensitivity under Ambiguity,” *American Economic Review*, 2024, *114* (12), 4091–4133.
- Jakubik, Adam and Roberta Piermartini**, “How WTO commitments tame uncertainty,” *European Economic Review*, 2023, *157*, 104495. Originally circulated as WTO Staff Working Paper ERSD-2019-06.
- Khan, Shafaat Yar and Armen Khederlarian**, “How Does Trade Respond to Anticipated Tariff Changes? Evidence from NAFTA,” *Journal of International Economics*, 2021, *133*, 103538.
- Klibanoff, Peter, Massimo Marinacci, and Sujoy Mukerji**, “A Smooth Model of Decision Making under Ambiguity,” *Econometrica*, 2005, *73* (6), 1849–1892.
- Kopecky, Karen A. and Richard M. H. Suen**, “Finite State Markov-Chain Approximations to Highly Persistent Processes,” *Review of Economic Dynamics*, 2010, *13* (3), 701–714.
- Maccheroni, Fabio, Massimo Marinacci, and Aldo Rustichini**, “Ambiguity Aversion, Robustness, and the Variational Representation of Preferences,” *Econometrica*, 2006, *74* (6), 1447–1498.
- Manova, Kalina, Dennis Novy, Thomas Sampson, and Aaron Tang**, “Tariff Confusion,” CEP Discussion Paper 2195, Centre for Economic Performance, London School of Economics 2026.

- Nishimura, Kiyohiko G. and Hiroyuki Ozaki**, “Irreversible investment and Knightian uncertainty,” *Journal of Economic Theory*, 2007, *136* (1), 668–694.
- Pierce, Justin R. and Peter K. Schott**, “The Surprisingly Swift Decline of US Manufacturing Employment,” *American Economic Review*, 2016, *106* (7), 1632–1662.
- Rouwenhorst, K. Geert**, “Asset Pricing Implications of Equilibrium Business Cycle Models,” in Thomas F. Cooley, ed., *Frontiers of Business Cycle Research*, Princeton, NJ: Princeton University Press, 1995, pp. 294–330.
- Simonovska, Ina and Michael E Waugh**, “The elasticity of trade: Estimates and evidence,” *Journal of International Economics*, 2014, *92* (1), 34–50.
- Soderbery, Anson**, “Trade elasticities, heterogeneity, and optimal tariffs,” *Journal of International Economics*, 2018, *114*, 44–62.
- Strzalecki, Tomasz**, “Axiomatic Foundations of Multiplier Preferences,” *Econometrica*, 2011, *79* (1), 47–73.
- Thijssen, Jacco J. J.**, “Incomplete markets, ambiguity, and irreversible investment,” *Journal of Economic Dynamics and Control*, 2011, *35* (6), 909–921.

Table 1: Ambiguity versus risk at the same entry threshold


	Parameter value	Elasticity of z_L^* with respect to τ_H
Baseline (standard risk)		0.38
<i>Threshold 5% higher via:</i>		
Higher likelihood, η	0.157	0.55
Higher persistence, ρ	0.930	0.55
Ambiguity, θ	0.297	0.76
<i>Threshold 10% higher via:</i>		
Higher likelihood, η	0.224	0.73
Higher persistence, ρ	0.998	0.73
Ambiguity, θ	0.168	1.15

Notes: Each row reports the elasticity of the entry threshold with respect to the tariff bound τ_H in an economy whose threshold is 5% or 10% higher than in a baseline standard-risk economy with $\eta = 0.1$, $\rho = 0.8$, and $\tau_H = 1.25$. The parameter value is the value of η or ρ under standard risk, or θ under ambiguity, that achieves this threshold increase, holding the other parameters at their baseline values. Ambiguity is modeled with the exact robust-control operator (14), not the quadratic approximation (16). Other parameters as in Figure 1.

Table 2: Assigned parameter values


Parameter	Meaning	Value
σ	Demand elasticity	3.17
r	Interest rate	0.04
δ_0	Productivity sensitivity of exit	21.04
δ_1	Asymptotic exit rate	0.02
ρ_ξ	Variable trade cost persistence	0.92
κ_0	Sunk entry cost	0.49
κ_1	Fixed continuation cost	0.28
ξ_H	High iceberg cost	4.34
ρ_z	Productivity shock persistence	0.65
σ_z	Productivity shock dispersion	1.47

Figure 1: Parameterizations where z_L^* is more sensitive under policy ambiguity to η and ρ



Notes: Shows regions where one, both, or neither of the right-hand sides of (30) and (31) are positive. Red: both; blue: only (31); green: neither. Solid and dashed curves show where the right-hand sides of (30) and (31), respectively, equal zero. In both panels, $\beta = 0.96$, $\delta = 0.9$, $\sigma = 4$, and $\tau_L = 1.0$. Uses the first-order approximation (20).

Figure 2: Estimated trade elasticity dynamics in model simulations



Notes: Figure shows estimates of trade elasticity dynamics from (41) for different levels of tariff persistence under robust-control ambiguity with $\theta \in \{\infty, 0.5, 0.1, 0.05, 0.01\}$. (a), (b), (c): Estimates for $\rho = 0.8$, $\rho = 0.95$, and $\rho = 0.99$, respectively. (d): Estimated long-run trade elasticity as a function of tariff persistence. X-axis: tariff persistence ρ . Y-axis: ratio of long-run trade elasticity estimate β_{14}^X to short-run trade elasticity β_0^X .

Online appendix

A Proofs

Throughout, write $\gamma \equiv \beta\delta \in (0, 1)$ and $\Delta := 1 - \gamma(\rho - \eta) > 0$. Let $T_s = D\tau_s^{1-\sigma}$ with $\sigma > 1$ so that $T_H < T_L$. Under SR (standard risk), the low-state threshold satisfies

$$z_L^* = \frac{\kappa}{\Phi}, \quad \Phi = \frac{\bar{T}}{1-\gamma} + \frac{T_L - \bar{T}}{\Delta}, \quad \bar{T} = p_H T_H + (1-p_H)T_L, \quad p_H = \frac{\eta}{\eta + 1 - \rho}.$$

Under RC (robust control), the quadratic representation gives

$$V_L^X(z) = b_L z - \frac{a_L}{2\theta} z^2, \quad \frac{a_L}{2\theta} (z_L^*)^2 - b_L z_L^* + \kappa = 0,$$

with $b_L > 0$, $a_L > 0$, and the closed form

$$z_L^{*,RC}(\theta) = \frac{b_L - \sqrt{b_L^2 - \frac{2a_L}{\theta} \kappa}}{a_L/\theta} = \frac{b_L - \sqrt{Q}}{a_L/\theta}.$$

All proofs using the quadratic approximation impose $\theta > 2a_L\kappa/b_L^2$, so $Q = b_L^2 - 2a_L\kappa/\theta > 0$. Proofs for the exact model are identified explicitly and do not impose this restriction.

A.1 Comparative statics w.r.t. κ

Claim (22). *For both setups, $\frac{\partial z_L^*}{\partial \kappa} > 0$.*

(SR). Since $z_L^* = \kappa/\Phi$ with $\Phi > 0$, we have $\partial_\kappa z_L^* = 1/\Phi > 0$. □

(RC). From the closed form, $\frac{\partial z_L^{*,RC}}{\partial \kappa} = \frac{1}{\sqrt{Q}} > 0$, since $Q > 0$. □

Claim (23). *Ambiguity increases the sensitivity to κ :*

$$\frac{\partial z_L^{*,RC}}{\partial \kappa} > \frac{\partial z_L^{*,SR}}{\partial \kappa}.$$

(RC vs. SR). $\partial_\kappa z_L^{*,RC} = 1/\sqrt{Q}$ and $\partial_\kappa z_L^{*,SR} = 1/b_L$. Since $\kappa > 0$, $0 < Q = b_L^2 - \frac{2a_L}{\theta}\kappa < b_L^2$, so $\sqrt{Q} < b_L$ and thus $1/\sqrt{Q} > 1/b_L$. □

Claim (24). *A higher ambiguity level strengthens the effect of κ :*

$$\frac{\partial^2 z_L^{*,RC}}{\partial \theta \partial \kappa} < 0.$$

(RC). Using $\partial_\kappa z_L^{*,\text{RC}} = Q^{-1/2}$ with $Q = b_L^2 - \frac{2a_L}{\theta}\kappa$,

$$\frac{\partial^2 z_L^{*,\text{RC}}}{\partial\theta\partial\kappa} = -\frac{1}{2}Q^{-3/2}\frac{\partial Q}{\partial\theta} = -\frac{a_L\kappa}{\theta^2 Q^{3/2}} < 0,$$

since $a_L, \kappa, \theta, Q > 0$. □

Claim (25). *Linearity (SR) and convexity (RC) in κ :*

$$\frac{\partial^2 z_L^{*,\text{SR}}}{\partial\kappa^2} = 0, \quad \frac{\partial^2 z_L^{*,\text{RC}}}{\partial\kappa^2} > 0.$$

(SR). $z_L^* = \kappa/\Phi$ implies $\partial_{\kappa\kappa}^2 z_L^* \equiv 0$. □

(RC). Differentiating $\partial_\kappa z_L^{*,\text{RC}} = Q^{-1/2}$ yields

$$\frac{\partial^2 z_L^{*,\text{RC}}}{\partial\kappa^2} = \frac{1}{2}Q^{-3/2}\left(\frac{2a_L}{\theta}\right) = \frac{a_L/\theta}{Q^{3/2}} > 0,$$

since $a_L, \theta, Q > 0$. □

Sunk-cost curvature in the exact RC model. Assume $0 < \eta, \rho < 1$, $T_L > T_H > 0$, and $\theta, \kappa > 0$. Write $V(z) = V^X(z)$ and let $J_{ss'} = P_{ss'}e^{-V_{s'}/\theta} / \sum_r P_{sr}e^{-V_r/\theta}$ be the tilted transition matrix, so that $\partial\mathcal{A}_s^{\text{RC}}/\partial V_{s'} = J_{ss'}$. Differentiating $V = zT + \gamma\mathcal{A}^{\text{RC}}[V]$ twice with respect to productivity gives

$$V_z = (I - \gamma J)^{-1}T, \quad V_{zz} = -\frac{\gamma}{\theta}(I - \gamma J)^{-1}w, \quad w_s = J_{sL}J_{sH}(V_{L,z} - V_{H,z})^2.$$

The inverse is nonnegative, $V_z > 0$, and

$$V_{L,z} - V_{H,z} = \frac{T_L - T_H}{1 - \gamma(J_{HH} - J_{LH})} > 0.$$

Since all entries of J are positive, $w > 0$ and $V_{zz} < 0$. The entry condition $V_L(z^*; \theta) = \kappa$ therefore implies

$$z_{\kappa\kappa}^* = -\frac{V_{L,zz}}{(V_{L,z})^3} > 0.$$

The robust-control operator satisfies $\mathcal{A}_s^{\text{RC}}[cV; c\theta] = c\mathcal{A}_s^{\text{RC}}[V; \theta]$ for $c > 0$, so if V is the value of exporting for productivity z and penalty θ , then cV is the value for cz and $c\theta$. Hence $z^*(c\kappa, c\theta) = cz^*(\kappa, \theta)$. Differentiating Euler's identity with respect to κ yields

$$\kappa z_{\kappa\kappa}^* + \theta z_{\kappa\theta}^* = 0, \quad z_{\kappa\theta}^* = -\frac{\kappa}{\theta} z_{\kappa\kappa}^* < 0.$$

Thus both signs hold for every finite $\theta > 0$, even though the threshold approaches a linear function of κ as $\theta \rightarrow 0$. $\square$

A.2 Comparative statics w.r.t. τ_H

Claim (26). *For each expectational setup (SR, RC) , the low-state entry threshold is increasing in the high tariff: $\partial z_L^*/\partial \tau_H > 0$.*

(SR). Write $z_L^* = \kappa/\Phi$ with

$$\Phi = \frac{\bar{T}}{1-\gamma} + \frac{T_L - \bar{T}}{\Delta}, \quad \bar{T} = p_H T_H + (1-p_H)T_L,$$

where $T_H = D\tau_H^{1-\sigma}$ and $\sigma > 1$ so that T_H is strictly *decreasing* in τ_H . Since Δ and p_H do not depend on τ_H ,

$$\frac{\partial \Phi}{\partial \tau_H} = \left(\frac{1}{1-\gamma} - \frac{1}{\Delta} \right) p_H \frac{\partial T_H}{\partial \tau_H} < 0,$$

because $\frac{1}{1-\gamma} - \frac{1}{\Delta} = \frac{\gamma(\eta+1-\rho)}{(1-\gamma)\Delta} > 0$ and $\partial T_H/\partial \tau_H < 0$. Hence $\partial_{\tau_H} z_L^* = -\kappa \Phi_{\tau_H}/\Phi^2 > 0$. $\square$

(RC). The low-state threshold satisfies $\kappa = b_L z - \frac{a_L}{2\theta} z^2$ with $b_L > 0$, $a_L > 0$. Differentiating implicitly w.r.t. τ_H ,

$$\left(b_L - \frac{a_L}{\theta} z \right) \frac{\partial z}{\partial \tau_H} = \frac{a_{L,\tau_H}}{2\theta} z^2 - b_{L,\tau_H} z.$$

Here $b_{L,\tau_H} < 0$ (because b_L is linear in T_H , which is decreasing in τ_H with a positive weight) and $a_{L,\tau_H} > 0$ (since $a_L \propto (b_H - b_L)^2$ and $|b_H - b_L|$ increases as T_H falls). The coefficient $b_L - \frac{a_L}{\theta} z = \sqrt{b_L^2 - \frac{2a_L}{\theta} \kappa} > 0$. Hence the right-hand side is positive and $\partial_{\tau_H} z_L^{*,RC} > 0$. $\square$

Claim (27). *Ambiguity strictly amplifies the τ_H -sensitivity:*

$$\frac{\partial z_L^{*,RC}}{\partial \tau_H} > \frac{\partial z_L^{*,SR}}{\partial \tau_H}.$$

(RC vs. SR). From the implicit formula,

$$\frac{\partial z_L^{*,RC}}{\partial \tau_H} = \frac{-b_{L,\tau_H} z}{\sqrt{Q}} + \frac{a_{L,\tau_H}}{2\theta} \frac{z^2}{\sqrt{Q}}, \quad Q := b_L^2 - \frac{2a_L}{\theta} \kappa.$$

By contrast, $\frac{\partial z_L^{*,SR}}{\partial \tau_H} = -\frac{\kappa b_{L,\tau_H}}{b_L^2}$. Subtracting and using $\kappa = b_L z - \frac{a_L}{2\theta} z^2$,

$$\frac{\partial z_L^{*,RC}}{\partial \tau_H} - \frac{\partial z_L^{*,SR}}{\partial \tau_H} = \left(-b_{L,\tau_H} \right) \left[\frac{z}{\sqrt{Q}} - \frac{\kappa}{b_L^2} \right] + \frac{a_{L,\tau_H}}{2\theta} \frac{z^2}{\sqrt{Q}}.$$

Since $b_{L,\tau_H} < 0$, $a_{L,\tau_H} > 0$, and $\sqrt{Q} < b_L$,

$$\frac{z}{\sqrt{Q}} - \frac{\kappa}{b_L^2} = z \left(\frac{1}{\sqrt{Q}} - \frac{1}{b_L} \right) + \frac{a_L z^2}{2\theta b_L^2} > 0,$$

so both terms on the right are strictly positive. Hence $\partial_{\tau_H} z_L^{*,RC} > \partial_{\tau_H} z_L^{*,SR}$. $\square$

Claim (28). *Cross-effects with the ambiguity parameter.*

$$\frac{\partial^2 z_L^{*,RC}}{\partial \theta \partial \tau_H} < 0.$$

(RC). Using the representation $\partial_{\tau_H} z = U/\sqrt{Q}$ with

$$U(\theta) = \frac{a_{L,\tau_H}}{2\theta} z^2 - b_{L,\tau_H} z \quad \text{and} \quad Q(\theta) = b_L^2 - \frac{2a_L}{\theta} \kappa,$$

we have

$$\frac{\partial}{\partial \theta} \left(\frac{\partial z}{\partial \tau_H} \right) = \frac{U_\theta}{\sqrt{Q}} - \frac{U Q_\theta}{2 Q^{3/2}}, \quad Q_\theta = \frac{2a_L \kappa}{\theta^2} > 0.$$

Since $a_{L,\tau_H} > 0$, $-b_{L,\tau_H} > 0$, and $z_\theta < 0$,

$$U_\theta = -\frac{a_{L,\tau_H}}{2\theta^2} z^2 + \frac{a_{L,\tau_H}}{\theta} z z_\theta - b_{L,\tau_H} z_\theta < 0.$$

Thus the first term is negative and the second term is also negative (because $U > 0$, $Q_\theta > 0$). Therefore $\partial^2 z_L^{*,RC} / (\partial \theta \partial \tau_H) < 0$. $\square$

A.3 Comparative statics w.r.t. η, ρ

Claim (29). *For $p \in \{\eta, \rho\}$ and for each expectational setup (SR, RC), the low-state entry threshold is increasing in p : $\partial z_L^* / \partial p > 0$.*

(SR). With $z_L^* = \kappa / \Phi$ and $\Phi = \bar{T} / (1 - \gamma) + (T_L - \bar{T}) / \Delta$ where $\bar{T} = p_H T_H + (1 - p_H) T_L$, $p_H = \eta / (\eta + 1 - \rho)$, and $\Delta = 1 - \gamma(\rho - \eta)$, we have

$$\Phi_p = \bar{T}_p \left(\frac{1}{1 - \gamma} - \frac{1}{\Delta} \right) - (T_L - \bar{T}) \frac{\Delta_p}{\Delta^2}.$$

Since $T_H < T_L$ (because $\sigma > 1$), $\bar{T}_\eta = (1 - \rho)(T_H - T_L) / (\eta + 1 - \rho)^2 < 0$ and $\bar{T}_\rho = \eta(T_H - T_L) / (\eta + 1 - \rho)^2 < 0$, while $\Delta_\eta = +\gamma$ and $\Delta_\rho = -\gamma$. Moreover $1/(1 - \gamma) - 1/\Delta = \gamma(\eta + 1 - \rho) / [(1 - \gamma)\Delta] > 0$. Hence $\Phi_p < 0$ for $p \in \{\eta, \rho\}$, and therefore

$$\frac{\partial z_L^*}{\partial p} = -\frac{\kappa \Phi_p}{\Phi^2} > 0.$$

□

(*RC, quadratic approximation*). Under the quadratic approximation, the low-state threshold solves $\frac{a_L}{2\theta}z^2 - b_Lz + \kappa = 0$ with $b_L > 0$, $a_L > 0$. Implicit differentiation gives

$$\frac{\partial z_L^{*,RC}}{\partial p} = \frac{(a_{L,p}/2\theta)z^2 - b_{L,p}z}{b_L - \frac{a_L}{\theta}z} = \frac{z}{\sqrt{b_L^2 - \frac{2a_L}{\theta}\kappa}} \left(\frac{a_{L,p}}{2\theta}z - b_{L,p} \right),$$

where $x_{L,p} \equiv \partial x_L / \partial p$. The derivatives of b_L are

$$b_{L,\eta} = \frac{\gamma(1-\gamma\rho)}{(1-\gamma)\Delta^2}(T_H - T_L) < 0, \quad b_{L,\rho} = \frac{\gamma^2\eta}{(1-\gamma)\Delta^2}(T_H - T_L) < 0,$$

since $T_H < T_L$. Hence $\partial_p z_L^{*,RC} > 0$ whenever $a_{L,p} \geq 0$. When $a_{L,p} < 0$, $\partial_p z_L^{*,RC} > 0$ if and only if

$$\frac{z_L^{*,RC}}{\theta} < \frac{2b_{L,p}}{a_{L,p}},$$

which holds for sufficiently large θ but can fail when ambiguity is strong. Since $a_L \propto (b_H - b_L)^2$ times the present value of the one-step switching variances, $a_{L,p} < 0$ means that an increase in p reduces the uncertainty the firm faces about its future profits, for example because it shrinks the gap between the values of the two states. Under the approximation, the resulting decline in the variance penalty can then outweigh the decline in expected profits. □

(*RC, exact*). The result holds for all $\theta > 0$ without the approximation. Fix $z > 0$ and write the exporter's value as the fixed point $V = \Gamma(V)$ of

$$\Gamma(V) := zT + \gamma \mathcal{A}^{RC}[V], \quad \mathcal{A}_s^{RC}[V] = -\theta \log \left[(1 - P_{sH})e^{-V_L/\theta} + P_{sH}e^{-V_H/\theta} \right],$$

with $T = (T_L, T_H)^\top$, $P_{LH} = \eta$, and $P_{HH} = \rho$. We proceed in four steps.

Step 1: Γ has a unique fixed point. $\mathcal{A}_s^{RC}$ is monotone ($V \leq W$ implies $\mathcal{A}_s^{RC}[V] \leq \mathcal{A}_s^{RC}[W]$) and translation invariant ($\mathcal{A}_s^{RC}[V + c\mathbf{1}] = \mathcal{A}_s^{RC}[V] + c$), so Γ satisfies Blackwell's sufficient conditions and is a contraction with modulus γ . Denote its fixed point by $V(z; P)$.

Step 2: $V_L > V_H$ at the fixed point. Each $\mathcal{A}_s^{RC}[V]$ is a certainty equivalent and therefore lies between $\min(V_L, V_H)$ and $\max(V_L, V_H)$, so $|\mathcal{A}_L^{RC}[V] - \mathcal{A}_H^{RC}[V]| \leq |V_L - V_H|$. Suppose $V_L \leq V_H$. Then

$$V_L - V_H = z(T_L - T_H) + \gamma(\mathcal{A}_L^{RC}[V] - \mathcal{A}_H^{RC}[V]) \geq z(T_L - T_H) - \gamma(V_H - V_L),$$

so $(1 - \gamma)(V_L - V_H) \geq z(T_L - T_H) > 0$, a contradiction.

Step 3: the value of exporting is decreasing in p . When $V_L > V_H$,

$$\frac{\partial \mathcal{A}_s^{RC}[V]}{\partial P_{sH}} = -\theta \frac{e^{-V_H/\theta} - e^{-V_L/\theta}}{(1 - P_{sH})e^{-V_L/\theta} + P_{sH}e^{-V_H/\theta}} < 0 :$$

shifting probability toward the state with the lower continuation value lowers the certainty equivalent. Let P' raise η (or ρ) relative to P , with associated operator Γ' , and let $V = V(z; P)$. By Step 2, $\Gamma'(V) \leq \Gamma(V) = V$, with strict inequality in the low state if η rises and in the high state if ρ rises. Since Γ' is monotone, iterating it from V produces a decreasing sequence that converges to its fixed point, so $V(z; P') \leq \Gamma'(V) \leq V$. If η rises, $V_L(z; P') \leq \Gamma'_L(V) < V_L$. If ρ rises, $V_H(z; P') < V_H$, and since $\mathcal{A}_L^{RC}$ is strictly increasing in V_H when $\eta > 0$, $V_L(z; P') = zT_L + \gamma \mathcal{A}_L^{RC}[V(z; P')] < zT_L + \gamma \mathcal{A}_L^{RC}[V] = V_L$. In both cases, $V_L(z; P') < V_L(z; P)$.

Step 4: the threshold is increasing in p . For $z' > z$, $\Gamma_{z'}(V(z; P)) = V(z; P) + (z' - z)T \geq V(z; P) + (z' - z)T_H \mathbf{1}$. Combining monotonicity and translation invariance, $V(z'; P) \geq V(z; P) + \frac{(z' - z)T_H}{1 - \gamma} \mathbf{1}$, so $V_L(z; P)$ is continuous, strictly increasing, and unbounded in z , and is zero at $z = 0$. The threshold $z_L^{*,RC}$, defined by $V_L(z_L^{*,RC}; P) = \kappa$ as in (9), is therefore unique. By Step 3, $V_L(z; P') < V_L(z; P)$ for every $z > 0$, so the threshold under P' is strictly higher. Hence $\partial_p z_L^{*,RC} > 0$ for all $\theta > 0$. $\square$

Claims (30)–(31). (RC vs. SR for the η - and ρ -partials.) Let $\hat{z}_L^{*,RC}$ denote the first-order approximation (20), let q , A , and B be defined as in the main text, and define

$$S_\eta := \frac{1}{\eta} - \frac{1 - \gamma\rho}{A} - \frac{3\gamma q}{B}, \quad S_\rho := \frac{3}{B} - \frac{2\rho - \eta}{A}.$$

Then

$$\text{sgn}\left(\frac{\partial \hat{z}_L^{*,RC}}{\partial \eta} - \frac{\partial z_L^{*,SR}}{\partial \eta}\right) = \text{sgn}(S_\eta), \quad \text{sgn}\left(\frac{\partial \hat{z}_L^{*,RC}}{\partial \rho} - \frac{\partial z_L^{*,SR}}{\partial \rho}\right) = \text{sgn}(S_\rho).$$

Proof. Assume $0 < \eta, \rho < 1$ and $0 < q < 1$, so that $A > 0$ and $B > 0$. Solving the two-state systems that define b_L and a_L gives

$$b_L = \frac{T_L B}{(1 - \gamma)\Delta}, \quad a_L = \frac{\gamma(T_L - T_H)^2 \eta A}{(1 - \gamma)\Delta^3}.$$

Since $z_L^{*,SR} = \kappa/b_L$, substituting into (20) shows that the Δ terms cancel from the ambiguity correction:

$$\hat{z}_L^{*,RC} - z_L^{*,SR} = \frac{\kappa^2 a_L}{2\theta b_L^3} = K \frac{\eta A}{B^3}, \quad K := \frac{\kappa^2 \gamma (1 - \gamma)^2 (1 - q)^2}{2\theta T_L} > 0.$$

K does not depend on η or ρ , so for $p \in \{\eta, \rho\}$,

$$\frac{\partial z_L^{*,\text{RC}}}{\partial p} - \frac{\partial z_L^{*,\text{SR}}}{\partial p} = K \frac{\eta A}{B^3} \frac{\partial}{\partial p} \log\left(\frac{\eta A}{B^3}\right).$$

Using $\partial_\eta A = -(1 - \gamma\rho)$, $\partial_\rho A = -\gamma(2\rho - \eta)$, $\partial_\eta B = \gamma q$, and $\partial_\rho B = -\gamma$,

$$\frac{\partial}{\partial \eta} \log\left(\frac{\eta A}{B^3}\right) = \frac{1}{\eta} - \frac{1 - \gamma\rho}{A} - \frac{3\gamma q}{B} = S_\eta, \quad \frac{\partial}{\partial \rho} \log\left(\frac{\eta A}{B^3}\right) = \frac{3\gamma}{B} - \frac{\gamma(2\rho - \eta)}{A} = \gamma S_\rho.$$

Since $K\eta A/B^3 > 0$ and $\gamma > 0$, the claimed equalities of signs follow. They are exact for the first-order approximation. The parameters θ and κ enter only through K , so they do not affect the signs. $\square$

Corollary (no “ η -only” region). *If the right-hand side in the η -condition of Claim (30) is positive, then so is the right-hand side in the ρ -condition of Claim (31).*

Proof. Using $B = (1 - \gamma\rho) + \gamma\eta q$,

$$\eta S_\eta - (1 - \gamma\rho) S_\rho = 1 - \frac{3[(1 - \gamma\rho) + \gamma\eta q]}{B} + \frac{(1 - \gamma\rho)[(2\rho - \eta) - \eta]}{A} = -2 + \frac{2(1 - \gamma\rho)(\rho - \eta)}{A}.$$

Since $A - (1 - \gamma\rho)(\rho - \eta) = (1 - \gamma\rho)(1 - \rho) + \gamma\rho(1 - \rho) = 1 - \rho$, this simplifies to

$$\eta S_\eta - (1 - \gamma\rho) S_\rho = -\frac{2(1 - \rho)}{A} < 0.$$

Because $\eta > 0$ and $1 - \gamma\rho > 0$, $S_\eta \geq 0$ implies $(1 - \gamma\rho)S_\rho > \eta S_\eta \geq 0$, so $S_\rho > 0$. Thus the pair $(S_\eta \geq 0, S_\rho < 0)$ cannot occur, and the zero contours of S_η and S_ρ cannot intersect in the open unit square. $\square$

Corollary (all three regions are nonempty). *For every $\gamma \in (0, 1)$ and $q \in (0, 1)$, each of the regions $\{S_\eta > 0, S_\rho > 0\}$, $\{S_\eta < 0, S_\rho > 0\}$, and $\{S_\eta < 0, S_\rho < 0\}$ has a nonempty interior in the (η, ρ) plane.*

Proof. By continuity, it suffices to find a limit point where the signs are strict in each case. *Both:* fix $\rho \in (0, 1)$ and let $\eta \rightarrow 0$. Then $S_\eta \rightarrow +\infty$, while $A \rightarrow 1 - \gamma\rho^2$ and $B \rightarrow 1 - \gamma\rho$, so

$$S_\rho \rightarrow \frac{3 - 2\rho - \gamma\rho^2}{(1 - \gamma\rho)(1 - \gamma\rho^2)} > 0.$$

ρ only: fix $\eta \in (1/2, 1)$ and let $\rho \rightarrow 0$. Then $A \rightarrow 1 - \eta$ and $B \rightarrow 1 + \gamma\eta q$, so

$$S_\eta \rightarrow \frac{1}{\eta} - \frac{1}{1 - \eta} - \frac{3\gamma q}{1 + \gamma\eta q} < 0, \quad S_\rho \rightarrow \frac{3}{1 + \gamma\eta q} + \frac{\eta}{1 - \eta} > 0.$$

Neither: let $\eta = \rho = t \rightarrow 1$. Then $A = 1 - t \rightarrow 0$ while $B = 1 - \gamma t(1 - q)$ stays bounded away from zero, so $S_\eta \rightarrow -\infty$ and $S_\rho \rightarrow -\infty$. $\square$

Corollary (tariff bound). *Holding γ , η , and ρ fixed, S_η and S_ρ are strictly decreasing in q . Hence, holding σ and τ_L fixed, a higher tariff bound τ_H expands the region where both are positive and contracts the region where both are negative.*

Proof. Differentiating,

$$\frac{\partial S_\eta}{\partial q} = -\frac{3\gamma(1 - \gamma\rho)}{B^2} < 0, \quad \frac{\partial S_\rho}{\partial q} = -\frac{3\gamma\eta}{B^2} < 0.$$

Since $q = T_H/T_L = (\tau_H/\tau_L)^{1-\sigma}$ with $\sigma > 1$, q is strictly decreasing in τ_H , so a higher tariff bound raises both S_η and S_ρ at every point in the open unit square. A point where both are positive therefore remains so, and a point where both are negative may cease to be so, but not conversely. $\square$

B Max-min multiple-priors ambiguity

B.1 Solution with max-min policy ambiguity

The [Gilboa and Schmeidler \(1989\)](#) approach to modeling ambiguity assumes that firms do not know the transition matrix P , but instead have a set of possible priors about that matrix and act as if they faced the worst possible element in that set. The one-step operator in this setup is

$$\mathcal{A}_s^{MM} [V_{s'}^X; \alpha] = \inf_{P' \in \mathcal{P}(\alpha)} \mathbb{E}_{P'} [V_{s'}^X] = \inf_{P' \in \mathcal{P}(\alpha)} \sum_{s'=L,H} P'_{s,s'} V_{s'}^X. \quad (\text{B.1})$$

The degree of ambiguity α controls the extent to which the set of priors $\mathcal{P}(\alpha)$ is allowed to diverge from the true matrix. I use a rectangular set of transition matrices, allowing each row to vary independently:

$$\mathcal{P}(\alpha) = \{P' : P'(s, \cdot) = (1 - \alpha)P(s, \cdot) + \alpha q_s, \quad q_s \in \text{conv}\{e_L, e_H\}, \quad s = L, H\}, \quad (\text{B.2})$$

where $\alpha \in [0, 1]$, $e_L = (1, 0)$ and $e_H = (0, 1)$ are point-mass probability vectors, and $\text{conv}\{e_L, e_H\}$ is the set of probability distributions over the two states. This implies that the firm simply puts the additional weight α on the worst-case outcome τ_H in both states, acting as if it faces the “tilted” transition matrix,

$$P'(\alpha) = \begin{bmatrix} 1 - \eta'(\alpha) & \eta'(\alpha) \\ 1 - \rho'(\alpha) & \rho'(\alpha) \end{bmatrix}, \quad (\text{B.3})$$

where $\eta'(\alpha) = \eta + \alpha(1 - \eta)$ and $\rho'(\alpha) = \rho + \alpha(1 - \rho)$.

The solution in this setup takes the same form as (13), with the elements of P replaced by the analogous elements of $P'(\alpha)$. The value of starting as an exporter in the low-tariff state is now

$$V_L^X(z_L^*(\alpha); \alpha) = z_L^*(\alpha) \left[\frac{\bar{T}(\alpha)}{1 - \beta\delta} + \frac{T_L - \bar{T}(\alpha)}{1 - \beta\delta(\rho'(\alpha) - \eta'(\alpha))} \right], \quad (\text{B.4})$$

where $\bar{T}(\alpha) = p'_H(\alpha)T_H + (1 - p'_H(\alpha))T_L$ is the unconditional average of T_s under the worst-case transition matrix $P'(\alpha)$. The entry threshold is

$$z_L^{*,MM} = \frac{\kappa}{\frac{\bar{T}(\alpha)}{1 - \beta\delta} + \frac{T_L - \bar{T}(\alpha)}{1 - \beta\delta(\rho'(\alpha) - \eta'(\alpha))}}. \quad (\text{B.5})$$

The threshold takes the same form as under standard policy risk but is increasing in the degree of ambiguity, α . When α increases, the expected present value of tariffs under the worst-case prior $P'(\alpha)$ increases, because the probability of switching to the high-tariff state, $\eta'(\alpha)$, and the probability of staying in that state, $\rho'(\alpha)$, both rise. When $\alpha = 0$, this setup

coincides with standard risk. When $\alpha = 1$, the firm acts as if it will face the high-tariff state in every period after the current one. Formally,

$$z_L^{*,MM} > z_L^{*,SR} \quad \forall \alpha > 0, \quad \frac{\partial z_L^{*,MM}}{\partial \alpha} > 0, \quad \forall \alpha \in [0, 1]. \quad (\text{B.6})$$

B.2 Comparative statics

This section states the comparative statics under max-min ambiguity together with their proofs. Notation follows Appendix A. In the MM case, the SR formula $z_L^* = \kappa/\Phi$ holds with the tilted transition matrix $P'(\alpha)$:

$$\eta' = \eta + \alpha(1 - \eta), \quad \rho' = \rho + \alpha(1 - \rho), \quad p'_H = \frac{\eta'}{\eta' + 1 - \rho'}, \quad \Delta'(\alpha) = 1 - \gamma(\rho' - \eta').$$

B.2.1 Sunk entry cost (κ)

As in the other two setups, the higher the sunk cost, the more productive a firm must be to enter:

$$\frac{\partial z_L^{*,MM}}{\partial \kappa} > 0 \quad \forall \alpha \in [0, 1]. \quad (\text{B.7})$$

Proof. Since $z_L^* = \kappa/\Phi$ with $\Phi > 0$, we have $\partial_\kappa z_L^* = 1/\Phi > 0$. $\square$

Because policy ambiguity reduces the certainty-equivalent of export profits, it makes the entry threshold more sensitive to the sunk cost:

$$\frac{\partial z_L^{*,MM}}{\partial \kappa} > \frac{\partial z_L^{*,SR}}{\partial \kappa} \quad \forall \alpha \in (0, 1]. \quad (\text{B.8})$$

(MM vs. SR). $\partial_\kappa z_L^{*,MM}(\alpha) = 1/\Phi(\alpha)$ and $\partial_\kappa z_L^{*,SR} = 1/\Phi(0)$. Because $z_L^{*,MM}(\alpha) = \kappa/\Phi(\alpha) > \kappa/\Phi(0) = z_L^{*,SR}$ for every $\alpha > 0$ (max-min tilt raises the threshold), it follows that $1/\Phi(\alpha) > 1/\Phi(0)$. Hence $\partial_\kappa z_L^{*,MM}(\alpha) > \partial_\kappa z_L^{*,SR}$. $\square$

A corollary is that a marginal increase in the sunk cost has a larger impact on exporting when ambiguity is stronger:

$$\frac{\partial^2 z_L^{*,MM}}{\partial \alpha \partial \kappa} > 0 \quad \forall \alpha \in (0, 1]. \quad (\text{B.9})$$

(MM). Since $z_L^{*,MM}(\alpha) = \kappa/\Phi(\alpha)$,

$$\frac{\partial^2 z_L^{*,MM}}{\partial \alpha \partial \kappa} = \frac{\partial}{\partial \alpha} \left(\frac{1}{\Phi(\alpha)} \right) = -\frac{\Phi_\alpha(\alpha)}{\Phi(\alpha)^2}.$$

But $\partial_\alpha z_L^{*,MM} = -\kappa \Phi_\alpha / \Phi^2 > 0$ for $\alpha \in [0, 1]$ (the threshold increases in α), hence $-\Phi_\alpha / \Phi^2 > 0$. $\square$

As under standard risk, the threshold increases linearly with the sunk cost:

$$\frac{\partial^2 z_L^{*,MM}}{\partial \kappa^2} = 0. \quad (\text{B.10})$$

Proof. $z_L^* = \kappa / \Phi$ implies $\partial_{\kappa\kappa}^2 z_L^* \equiv 0$. $\square$

This is because the firm's objective in the max-min setup is still linear in the payoff stream, just as in the standard risk setup, whereas robust control penalizes the variance in future payoffs.

B.2.2 High tariff value (τ_H)

As in the other two setups, the more tariffs can rise, the more productive a firm must be to make entry profitable:

$$\frac{\partial z_L^{*,MM}}{\partial \tau_H} > 0 \quad \forall \alpha \in [0, 1]. \quad (\text{B.11})$$

(MM). The same calculation as in the SR proof of (26) in Appendix A holds with p_H and Δ replaced by the worst-case objects $p'_H(\alpha)$ and $\Delta'(\alpha)$, which also do not depend on τ_H . Therefore $\Phi_{\tau_H}(\alpha) < 0$ and $\partial_{\tau_H} z_L^{*,MM} > 0$. $\square$

Trade-policy ambiguity makes the potential increase in tariffs more important:

$$\frac{\partial z_L^{*,MM}}{\partial \tau_H} > \frac{\partial z_L^{*,SR}}{\partial \tau_H} \quad \forall \alpha \in (0, 1]. \quad (\text{B.12})$$

In the max-min setup, this is because there is more weight on the high-tariff state in the worst-case transition matrix $P'(\alpha)$.

(MM vs. SR). Using the SR/MM formulas,

$$\frac{\partial z_L^*}{\partial \tau_H} = \frac{\kappa}{\Phi(\cdot)^2} \left(\frac{1}{1-\gamma} - \frac{1}{\Delta(\cdot)} \right) p_H(\cdot) \left| \frac{\partial T_H}{\partial \tau_H} \right|,$$

where $(\cdot)$ denotes either the baseline (p_H, Δ) or the worst-case (p'_H, Δ') . Using $[1/(1-\gamma) - 1/\Delta]p_H = \gamma\eta/[(1-\gamma)\Delta]$ and its tilted counterpart gives

$$\frac{\partial_{\tau_H} z_L^{*,MM}}{\partial_{\tau_H} z_L^{*,SR}} = \left(\frac{\Phi}{\Phi(\alpha)} \right)^2 \frac{\eta'}{\eta} \frac{\Delta}{\Delta'}.$$

First, $\Phi(\alpha) < \Phi$ for $\alpha > 0$, so the prefactor $(\Phi/\Phi(\alpha))^2 > 1$. Second,

$$\frac{\eta'}{\eta} \cdot \frac{\Delta}{\Delta'} \geq 1,$$

because $\eta' = \eta + \alpha(1 - \eta)$ and $\Delta' = 1 - \gamma(\rho' - \eta') = 1 - \gamma(1 - \alpha)(\rho - \eta)$ imply

$$\frac{d}{d\alpha} \log\left(\frac{\eta'}{\Delta'}\right) = \frac{1 - \eta}{\eta'} - \frac{\gamma(\rho - \eta)}{\Delta'} \geq 0 \iff 1 - \eta - \gamma(\rho - \eta) \geq 0,$$

which holds since $\gamma \in (0, 1)$ and $\rho, \eta \in [0, 1]$ (indeed $1 - \eta - \gamma(\rho - \eta) = 1 - \gamma\rho - (1 - \gamma)\eta \geq 0$). Therefore the product above is ≥ 1 with strict $>$ for $\alpha > 0$. Hence $\partial_{\tau_H} z_L^{*,MM} > \partial_{\tau_H} z_L^{*,SR}$. $\square$

As with the sunk cost, changes in the value of the high tariff have more impact on exporting when the degree of ambiguity is greater:

$$\frac{\partial^2 z_L^{*,MM}}{\partial \alpha \partial \tau_H} > 0 \quad \forall \alpha \in (0, 1]. \quad (\text{B.13})$$

(MM). As shown above,

$$\frac{\partial z_L^{*,MM}}{\partial \tau_H} = \frac{\kappa \gamma}{1 - \gamma} \frac{\eta'}{\Delta' \Phi(\alpha)^2} \left| \frac{\partial T_H}{\partial \tau_H} \right|.$$

The factor η'/Δ' is weakly increasing in α (same inequality as in the proof of (B.12)), while $\Phi(\alpha)$ is strictly decreasing in α ; hence the whole expression is strictly increasing in α on $[0, 1)$. $\square$

B.2.3 Likelihood and persistence of high tariffs (η, ρ)

As in the other two setups, the more likely tariffs are to rise and the more persistent such an increase is, the more productive a firm must be to enter when tariffs are low, except when $\alpha = 1$:

$$\frac{\partial z_L^{*,MM}}{\partial p} > 0 \quad \forall \alpha \in [0, 1), \quad p = \eta, \rho. \quad (\text{B.14})$$

(MM). Under rectangular max-min, $z_L^{*,MM}(\alpha) = \kappa/\Phi(\alpha)$ with the same $\Phi(\cdot)$ but evaluated at the tilted matrix $P'(\alpha)$ and, crucially,

$$\Phi_p(\alpha) = (1 - \alpha) \mathcal{B}_p(\eta'(\alpha), \rho'(\alpha)),$$

where $\mathcal{B}_p$ is the SR bracket Φ_p from the proof of (29) in Appendix A. Since $\mathcal{B}_p < 0$ at any admissible (η, ρ) and $(1 - \alpha) \geq 0$, it follows that $\Phi_p(\alpha) < 0$ for $\alpha \in [0, 1)$ and $\partial_p z_L^{*,MM}(\alpha) > 0$. As $\alpha \uparrow 1$, $(1 - \alpha) \rightarrow 0$ and $\partial_p z_L^{*,MM}(\alpha) \rightarrow 0$. (SR/MM objects as defined in Section 3.3 and Appendix B.1.) $\square$

In the max-min setup, there exists a threshold value $\bar{\alpha}_p$ for each transition probability $p = \eta, \rho$ such that $z_L^{*,MM}$ is less sensitive than $z_L^{*,SR}$ to changes in that probability for $\alpha > \bar{\alpha}_p$. The intuition is that as α goes to one, the firm acts as if it will always face the high-tariff state in the future regardless of η and ρ and does not react to changes in these parameters. Below the threshold $\bar{\alpha}_p$, z_L^* is more sensitive to changes in p than under standard risk. Formally,

$$\lim_{\alpha \rightarrow 1} \frac{\partial z_L^{*,MM}}{\partial p} = 0, \quad \exists \bar{\alpha}_p \in [0, 1) \text{ sign} \left(\frac{\partial z_L^{*,MM}}{\partial p} - \frac{\partial z_L^{*,SR}}{\partial p} \right) = \text{sign}(\bar{\alpha}_p - \alpha), \quad p = \eta, \rho \quad (\text{B.15})$$

For $p = \eta$, $\bar{\alpha}_\eta = 0$ in all parameterizations: $z_L^{*,MM}$ is always less sensitive than $z_L^{*,SR}$ to changes in η under max-min ambiguity, and $\partial z_L^{*,MM}/\partial p$ converges monotonically downward to zero as α goes to one. For $p = \rho$, on the other hand, there are some parameterizations where z_L^* is always less sensitive to ρ than under standard risk (i.e., $\bar{\alpha}_\rho = 0$), and there are others where it is more sensitive to ρ for low levels of ambiguity and less sensitive for higher levels (i.e., $\bar{\alpha}_\rho > 0$).

Figure B.1 illustrates this using a numerical example. Panel (a) plots the ratio of $\partial z_L^{*,MM}/\partial \eta$ to $\partial z_L^{*,SR}/\partial \eta$ as a function of α for several different values of η . In all cases, this ratio equals one for $\alpha = 0$ and declines as α rises. Panel (b) plots the ratio of $\partial z_L^{*,MM}/\partial \rho$ to $\partial z_L^{*,SR}/\partial \rho$. When η is low (i.e., high tariffs are unlikely), the entry threshold is more sensitive to ρ under max-min ambiguity than standard risk unless α is very high. When η is sufficiently high, the entry threshold is less sensitive to ρ under max-min ambiguity for all values of α . In general, varying ρ while holding fixed η changes the location of $\bar{\alpha}_\rho$, but does not change whether $\bar{\alpha}_\rho > 0$.

Proof. From the expression $\partial_p z_L^{*,MM}(\alpha) = -\kappa \Phi_p(\alpha)/\Phi(\alpha)^2$ with $\Phi_p(\alpha) = (1-\alpha)\mathcal{B}_p(\eta'(\alpha), \rho'(\alpha))$ and $\mathcal{B}_p < 0$, we have $\partial_p z_L^{*,MM}(\alpha) > 0$ for $\alpha \in [0, 1)$ and the limit to 0 as $\alpha \uparrow 1$.


Comparing to SR, define the continuous ratio

$$R_p(\alpha) := \frac{\partial_p z_L^{*,MM}(\alpha)}{\partial_p z_L^{*,SR}} = (1-\alpha) \left(\frac{\Phi(0)}{\Phi(\alpha)} \right)^2 \frac{|\mathcal{B}_p(\eta'(\alpha), \rho'(\alpha))|}{|\mathcal{B}_p(\eta, \rho)|}.$$

We have $R_p(0) = 1$ and $\lim_{\alpha \rightarrow 1} R_p(\alpha) = 0$, so there exists $\bar{\alpha}_p \in [0, 1)$ at which $R_p(\bar{\alpha}_p) = 1$. For $p = \eta$ one can show $R'_\eta(\alpha) < 0$ on $(0, 1)$: $1 - \alpha$ declines, $\Phi(\alpha)$ declines, and $|\mathcal{B}_\eta(\eta'(\alpha), \rho'(\alpha))|$ also declines because $\eta'(\alpha)$ enters $\bar{T}$ and Δ' so as to reduce $|\Phi_\eta|$. Hence $R_\eta(\alpha) < 1$ for all $\alpha > 0$ and $\bar{\alpha}_\eta = 0$. For $p = \rho$, $|\mathcal{B}_\rho(\eta'(\alpha), \rho'(\alpha))|$ may initially *rise* with α (the persistence channel dominates) before the $(1 - \alpha)$ dampening takes over, yielding either $\bar{\alpha}_\rho = 0$ or an interior $\bar{\alpha}_\rho \in (0, 1)$ depending on primitives—exactly the patterns in Figure B.1. $\square$

In the max-min setup, what matters is the degree of ambiguity; trade is more sensitive

Figure B.1: Sensitivity of z_L^* under max-min ambiguity to transition probabilities



Notes: (a) Partial derivative of z_L^* with respect to η under max-min ambiguity relative to the same partial derivative under standard risk, as a function of the degree of ambiguity α . (b) Same as (a), but with respect to ρ . In both panels, $\beta = 0.96$, $\delta = 0.9$, $\sigma = 4$, $\tau_L = 1.0$, and $\tau_H = 1.25$.

to increases in the persistence of high tariffs at low levels of ambiguity and less sensitive for higher levels.

B.2.4 Idiosyncratic productivity (z)

Max-min and robust control ambiguity have similar effects on export participation and the manner in which it depends on the model's parameters. However, there is one way in which the two approaches to modeling trade-policy ambiguity work differently: max-min ambiguity affects all firms in the same way, but high-productivity firms are more sensitive to ambiguity under robust control.

To see this, first note that the one-step operator $\mathcal{A}_s$ under robust control can be written as

$$\mathcal{A}_s^{RC} [V_{s'}^X; \theta] = \sum_{s'=L,H} \tilde{P}_{s,s'}(z; \theta) V_{s'}^X + \theta \sum_{s'=L,H} \tilde{P}_{s,s'}(z; \theta) \log \left[\frac{\tilde{P}_{s,s'}(z; \theta)}{P_{s,s'}} \right] \quad (\text{B.16})$$

with “tilted” probabilities given by

$$\tilde{P}_{s,s'}(z; \theta) = \frac{P_{s,s'} e^{-V_{s'}^X(z)/\theta}}{\sum_{r=L,H} P_{s,r} e^{-V_r^X(z)/\theta}} \quad (\text{B.17})$$

that depend on productivity z as well as the degree of ambiguity θ .

Proof. The Gibbs variational formula states that for any probability vector p and any vector

v ,

$$-\theta \log \sum_i p_i e^{-v_i/\theta} = \min_{q \in \Delta} \left\{ \sum_i q_i v_i + \theta \sum_i q_i \log \frac{q_i}{p_i} \right\}.$$

In our case, the Lagrangian for the minimization is

$$\mathcal{L}(q, \lambda) = \sum_{s'} q_{s'} V_{s'}^X(z) + \theta \sum_{s'} q_{s'} \log \frac{q_{s'}}{P_{ss'}} + \lambda \left(\sum_{s'} q_{s'} - 1 \right).$$

First-order conditions give $V_{s'}^X(z) + \theta(\log q_{s'} - \log P_{ss'} + 1) + \lambda = 0$, hence $q_{s'} \propto P_{ss'} e^{-V_{s'}^X(z)/\theta}$; normalizing $\sum q_s = 1$ delivers the desired solution for $\tilde{P}$. Substituting $q = \tilde{P}_{\cdot}(z; \theta)$ into the entropy representation gives the desired expression for $\mathcal{A}^{RC}$. $\square$

We can write the tilted odds ratio of the high-tariff state as

$$\frac{\tilde{P}_{s,H}(z; \theta)}{\tilde{P}_{s,L}(z; \theta)} = \frac{P_{s,H}}{P_{s,L}} e^{[V_L^X(z) - V_H^X(z)]/\theta} := \frac{P_{s,H}}{P_{s,L}} e^{\Delta V^X(z)/\theta} \quad (\text{B.18})$$

The value gap $\Delta V^X(z)$ is increasing in z , and so the odds ratio is as well:

$$\frac{\partial}{\partial z} \log \left[\frac{\tilde{P}_{s,H}(z; \theta)}{\tilde{P}_{s,L}(z; \theta)} \right] > 0. \quad (\text{B.19})$$

Proof. Taking logs of (B.18) and differentiating gives $\frac{\partial}{\partial z} \log \frac{\tilde{P}_{s,H}(z; \theta)}{\tilde{P}_{s,L}(z; \theta)} = \frac{1}{\theta} \frac{d}{dz} (V_L^X(z) - V_H^X(z))$. It remains to show $\Delta'(z) := \frac{d}{dz} (V_L^X - V_H^X) > 0$. Exporter values solve

$$V^X(z) = zT + \gamma \mathcal{A}^{RC}[V^X(z)], \quad T = \begin{bmatrix} T_L \\ T_H \end{bmatrix}, \quad \gamma = \beta\delta \in (0, 1).$$

Differentiate w.r.t. z :

$$(I - \gamma J(z)) V^{X'}(z) = T, \quad J_{ss'}(z) := \frac{\partial \mathcal{A}_s^{RC}}{\partial V_{s'}^X} [V^X(z)].$$

The representation in (B.16) gives $J_{ss'}(z) = \tilde{P}_{ss'}(z; \theta)$. Write the off-diagonals as $b := J_{LH}(z)$ and $c := J_{HL}(z)$, so $J = \begin{bmatrix} 1-b & b \\ c & 1-c \end{bmatrix}$. Then $V' = (V'_L, V'_H)^\top$ solves

$$\begin{aligned} (1 - \gamma(1 - b))V'_L - \gamma b V'_H &= T_L, \\ -\gamma c V'_L + (1 - \gamma(1 - c))V'_H &= T_H. \end{aligned}$$

Subtracting the two equations and solving for $\Delta'(z) = V'_L - V'_H$ yields

$$\Delta'(z) = \frac{(1 - \gamma)(T_L - T_H)}{(1 - \gamma)^2 + \gamma(1 - \gamma)(b + c)} = \frac{T_L - T_H}{(1 - \gamma) + \gamma(b + c)} > 0,$$

because $T_L > T_H$ and $b, c \in [0, 1]$. Hence $\Delta'(z) > 0$ and (B.19) follows. $\square$

Thus, high-productivity firms put relatively more weight on the high-tariff state than low-productivity firms under robust control. By contrast, under max-min ambiguity, all firms put the same additional weight α on the high-tariff state.


In the simple model I have analyzed in Section 3 the entry decision is the only margin of interest so there is no scope for this difference to affect any other outcomes, such as exit, firm growth, etc. In a quantitative model with ingredients like endogenous exit and exporter life cycles, this difference becomes more important. I take that up in Section 4.

B.3 Quantitative results

Figure B.2 shows the results of the simulation exercise described in Section 4 for max-min ambiguity. When tariffs are less persistent ($\rho = 0.8$), the long-run response of trade is similar regardless of the degree of ambiguity; even under standard policy risk, firms do not respond much to transitory reforms because they know they are likely to be reversed as in Alessandria et al. (2025a). When tariffs are more persistent ($\rho = 0.95$ and $\rho = 0.99$), ambiguity attenuates the long-run trade response to a greater degree, because firms act as if tariff reductions are more likely to be reversed than they really are. When ambiguity is severe ($\alpha = 0.75$), the long-run trade response is similar regardless of the true level of tariff persistence because firms always act as if a switch from low tariffs to high tariffs is likely to occur and likely to last.

Figure B.2(d) summarizes these results by plotting the ratio of the long-run trade elasticity estimate, β_{14}^X , to the short-run estimate, β_0^X .

Figure B.2: Estimated trade elasticity dynamics in model simulations: max-min ambiguity



Notes: Figure shows estimates of trade elasticity dynamics from (41) for different levels of tariff persistence under max-min ambiguity with $\alpha \in \{0, 0.1, 0.25, 0.5, 0.75\}$. (a), (b), (c): Estimates for $\rho = 0.8$, $\rho = 0.95$, and $\rho = 0.99$, respectively. (d): Estimated long-run trade elasticity as a function of tariff persistence. X-axis: tariff persistence ρ . Y-axis: ratio of long-run trade elasticity estimate β_{14}^X to short-run trade elasticity β_0^X .

C Accuracy of the quadratic approximation

The comparative statics in Section 3 use the quadratic approximation (16), which follows from a second-order expansion of the robust-control operator (14) in $1/\theta$. When continuation values are normally distributed, this expansion is not an approximation at all—it is an alternative but exact representation. When tariffs follow a two-state Markov process, however, higher moments of the distribution of continuation values matter, and the approximation is accurate only when θ is large. This appendix assesses how large θ must be by solving the exact, unapproximated model numerically.

In the exact model, the value of exporting solves $V_s^X(z) = zT_s + \delta\beta\mathcal{A}_s^{RC} [V_{s'}^X(z); \theta]$, and the low-tariff threshold solves $V_L^X(z_L^*) = \kappa$. I solve these three equations for (V_L^X, V_H^X, z_L^*) using Newton’s method, compute first derivatives of the threshold using the implicit function theorem, compute the sunk-cost second derivatives analytically as shown in Appendix A, and compute the remaining second derivatives using numerical differentiation. Because z_L^*/κ depends only on θ/κ , I normalize $\kappa = 1$. As in Figure 1, $\beta = 0.96$, $\delta = 0.9$, $\sigma = 4$, and $\tau_L = 1$. I consider $\tau_H \in \{1.05, 1.25, 1.50\}$ and $\eta, \rho \in [0.02, 0.98]$, and I measure the degree of ambiguity relative to the bound $\theta_{\min} \equiv 2a_L\kappa/b_L^2$ below which the quadratic approximation has no solution. The exact model has a solution for every $\theta > 0$, and as θ goes to zero its threshold converges to the max-min threshold with $\alpha = 1$.

Table C.1 compares the thresholds. Relative to standard risk, the quadratic and first-order thresholds depend only on $\theta/\theta_{\min}$; the exact threshold also depends on η , ρ , and τ_H . When ambiguity is moderate (θ at least ten times $\theta_{\min}$), both approximations are accurate. As ambiguity becomes stronger, they become less accurate, especially when tariff hikes are rare ($\eta < 0.1$). In this case, the distribution of continuation values is highly skewed, which a second-order expansion ignores, and the approximations understate the effect of ambiguity on the entry threshold. Near the bound, the quadratic threshold is twice the standard-risk threshold, which overstates the effect of ambiguity in most parameterizations.

Table C.2 reports the share of parameterizations in which each comparative static in Section 3 holds in the exact model. All hold throughout the grid, including when $\theta < \theta_{\min}$, except the first-order sign conditions for the hazard and persistence of high tariffs. In particular, the sunk-cost convexity and cross-partial results hold for every finite $\theta > 0$, as proved in Appendix A. Under strong ambiguity their magnitudes can be extremely small, so finite differences can incorrectly report zero or reversed signs.

The sign conditions (30)–(31) use the first-order approximation, so they describe how a small degree of ambiguity alters the threshold’s sensitivity to η and ρ . In the exact model, these signs match the first-order conditions in nearly every parameterization when ambiguity is weak, but less often as it becomes stronger. Figure C.1 shows the results. The three-region structure survives: in no parameterization is the threshold more sensitive to η

Table C.1: Entry threshold under ambiguity relative to standard risk

$\theta/\theta_{\min}$	Quadratic	First order	Exact model	
			$\eta \geq 0.1$	$\eta < 0.1$
100	1.003	1.003	1.003 [1.002, 1.003]	1.003 [1.003, 1.003]
10	1.026	1.025	1.026 [1.020, 1.031]	1.052 [1.031, 1.119]
3	1.101	1.083	1.085 [1.038, 1.128]	1.223 [1.096, 1.647]
1.5	1.268	1.167	1.156 [1.043, 1.298]	1.449 [1.107, 2.043]
1	2.000	1.250	1.203 [1.045, 1.446]	1.535 [1.111, 2.179]

Notes: Entries are $z_L^{*,RC}/z_L^{*,SR}$. $\theta_{\min} = 2a_L\kappa/b_L^2$ is the smallest θ for which the quadratic approximation (16) has a solution. Quadratic: root of the quadratic approximation. First order: (20). Exact model: median and, in brackets, 10th and 90th percentiles across $\tau_H \in \{1.05, 1.25, 1.50\}$ and $\eta, \rho \in [0.02, 0.98]$. Other parameters as in Figure 1.

Table C.2: Share of parameterizations where comparative statics hold in the exact model

Result	$\theta/\theta_{\min}$					
	100	10	3	1	0.3	0.1
$z_L^{*,RC} > z_L^{*,SR}$, $\partial z_L^{*,RC}/\partial\theta < 0$	1.00	1.00	1.00	1.00	1.00	1.00
(22)–(23), (26)–(29)	1.00	1.00	1.00	1.00	1.00	1.00
(24)–(25)	1.00	1.00	1.00	1.00	1.00	1.00
(30): sign matches first order	0.99	0.91	0.82	0.70	0.68	0.68
(31): sign matches first order	1.00	0.94	0.82	0.62	0.35	0.18

Notes: Share of parameterizations in which each result holds in the exact model, across $\tau_H \in \{1.05, 1.25, 1.50\}$ and $\eta, \rho \in [0.02, 0.98]$. $\theta_{\min}$ is defined as in Table C.1; for $\theta < \theta_{\min}$, only the exact model has a solution. Other parameters as in Figure 1.


but less sensitive to ρ than under standard risk. However, as ambiguity becomes stronger, the region where the threshold is more sensitive to both parameters shrinks and the region where it is less sensitive to both expands. This mirrors the max-min setup in Appendix B: as ambiguity becomes severe, the firm acts as if high tariffs are inevitable, and the threshold becomes insensitive to the likelihood and persistence of tariff increases.

Section 3.6 compares ambiguity and risk at the same entry threshold. Both of its results hold exactly in the robust-control model, for any $\theta > 0$.

The sunk-cost result (32) follows from the homogeneity property $z_L^{*,RC}(c\kappa, c\theta) = c z_L^{*,RC}(\kappa, \theta)$ shown in Appendix A. By Euler's theorem, $\kappa \partial z_L^{*,RC}/\partial\kappa + \theta \partial z_L^{*,RC}/\partial\theta = z_L^{*,RC}$, which gives the equality in (32). The inequality holds because the threshold is decreasing in θ : the derivative of the operator with respect to θ equals the relative entropy of the tilted distribution (B.17) with respect to P , which is positive whenever $V_L^X \neq V_H^X$, so the value of exporting is increasing in θ by the same monotonicity argument as in the proof of (29).

The bound result rests on a property of standard risk. The value of exporting for a firm

Figure C.1: Parameterizations where z_L^* is more sensitive under policy ambiguity to η and ρ : exact model



Notes: Red: z_L^* is more sensitive to both η and ρ under policy ambiguity than under standard risk; blue: more sensitive only to ρ ; green: less sensitive to both. Solid and dashed curves show where the differences in sensitivity to η and ρ , respectively, equal zero. Panel (a) uses the first-order conditions (30)–(31); panels (b)–(d) use the exact model with decreasing values of θ (increasing degrees of ambiguity). In all panels, $\beta = 0.96$, $\delta = 0.9$, $\sigma = 4$, $\tau_L = 1.0$, and $\tau_H = 1.25$.

with productivity one is $b_L = c_0 T_L + C T_H$, where c_0 and C are the elements of the first row of $(I - \gamma P)^{-1}$, that is, the expected discounted time spent in the low- and high-tariff states. These sum to $1/(1 - \gamma)$, so $b_L = T_L/(1 - \gamma) - C(T_L - T_H)$. Since C does not depend on τ_H , the elasticity of the threshold $z_L^{*,SR} = \kappa/b_L$ with respect to the bound is

$$\frac{\partial \log z_L^{*,SR}}{\partial \log \tau_H} = f(C) \equiv -\frac{\tau_H C}{T_L/(1 - \gamma) - C(T_L - T_H)} \frac{\partial T_H}{\partial \tau_H}, \quad (\text{C.1})$$

which is increasing in C . Given the tariffs, the threshold pins down

$$C = \frac{T_L/(1 - \gamma) - \kappa/z_L^{*,SR}}{T_L - T_H},$$

so every specification of η and ρ with the same threshold has the same bound elasticity.

Under robust control, the bound elasticity takes the same form. Differentiating the Bellman equation $V^X(z) = zT + \gamma \mathcal{A}^{RC}[V^X(z)]$ and using the fact that the derivative of $\mathcal{A}_s^{RC}$ with respect to $V_{s'}^X$ is the tilted probability $\tilde{P}_{s,s'}$ (B.17), the implicit function theorem gives $\partial \log z_L^{*,RC} / \partial \log \tau_H = f(\tilde{C})$, where $\tilde{C}$ is the (L, H) element of $(I - \gamma \tilde{P})^{-1}$, the expected discounted time spent in the high-tariff state under the tilted transition matrix. Moreover, $\tilde{b}_L \equiv T_L/(1 - \gamma) - \tilde{C}(T_L - T_H)$ equals $\partial V_L^X / \partial z$, the marginal value of productivity at the threshold. The robust-control operator is concave and the Bellman equation is linear in z , so by the same monotone-iteration argument as in the proof of (29), the value of exporting is concave in z . It is strictly concave in the low-tariff state because $V_L^X - V_H^X$ is strictly increasing in z (B.19), and the operator is strictly concave in every direction other than a constant shift. Since $V_L^X(0) = 0$, strict concavity implies that the marginal value of productivity is below its average value: $\tilde{b}_L < V_L^X(z_L^{*,RC})/z_L^{*,RC} = \kappa/z_L^{*,RC}$. The right-hand side is the value b_L that pins down C in any specification of risk with the same threshold. Hence $\tilde{C} > C$, and since f is increasing, the bound elasticity under ambiguity exceeds the bound elasticity under any specification of risk with the same threshold. Intuitively, more productive firms have more at stake and put more weight on the high-tariff state, so the marginal entrant is more exposed to the tariff bound than its average value of exporting reveals.

Table C.3 quantifies these differences. For each parameterization in the grid used above, it reports the ratio of the bound elasticity under ambiguity to the bound elasticity under standard risk with the same threshold, computed from (C.1), and the sunk-cost elasticity (32). The ratio is largest at moderate degrees of ambiguity. It approaches one as θ becomes large, because the economy approaches standard risk, and as θ goes to zero, because the firm then acts as if it will face the high tariff in every period after the current one, which is itself a specification of risk.

D AR(1) tariff process

The quantitative model in Section 4 assumes that tariffs follow a symmetric two-state Markov process: the tariff is either $\tau_L = 1$ or $\tau_H = 1.3$, and it stays in its current state with probability ρ . A natural question is whether the results depend on this assumption. To answer it, I replace the two-state process with a discretized AR(1) process and repeat the simulation exercise.

Table C.3: Ambiguity versus risk at the same entry threshold: all parameterizations

$\theta/\theta_{\min}$	Bound elasticity relative to risk		Sunk-cost elasticity	
	Median	Range	Median	Range
100	1.02	[1.00, 2.75]	1.00	[1.00, 1.02]
10	1.13	[1.00, 4.32]	1.03	[1.00, 1.75]
3	1.24	[1.00, 3.08]	1.08	[1.00, 1.62]
1	1.17	[1.00, 2.38]	1.07	[1.00, 1.52]
0.3	1.05	[1.00, 1.30]	1.02	[1.00, 1.16]
0.1	1.02	[1.00, 1.09]	1.01	[1.00, 1.05]

Notes: Bound elasticity relative to risk: ratio of the elasticity of z_L^* with respect to τ_H in the exact robust-control model to the same elasticity under standard risk with the same threshold, computed from (C.1). Sunk-cost elasticity: (32). Median and range across $\tau_H \in \{1.05, 1.25, 1.50\}$ and $\eta, \rho \in [0.02, 0.98]$. $\theta_{\min}$ is defined as in Table C.1. Other parameters as in Figure 1.

I discretize a symmetric AR(1) process for the tariff using the method of Rouwenhorst (1995), which Kopecky and Suen (2010) show is accurate for highly persistent processes. The tariff takes N values evenly spaced between $\tau_L = 1$ and τ_H . With $N = 2$, this nests exactly the process used in Section 4. For any N , the first-order autocorrelation of tariffs is $2\rho - 1$, the mean tariff is $(\tau_L + \tau_H)/2$, and the standard deviation of tariffs is $(\tau_H - \tau_L)/(2\sqrt{N-1})$. Thus, adding states while holding the bound fixed makes tariffs less volatile. I consider two families of processes with $N \in \{3, 5, 9\}$: one that holds the bound fixed at $\tau_H = 1.3$ (“fixed bound”), and one that raises the bound to $\tau_H = 1 + 0.3\sqrt{N-1}$ to hold the standard deviation fixed at its baseline value of 0.15 (“fixed variance”). In the second family, the mean tariff rises with N because tariffs cannot fall below $\tau_L = 1$. Everything else is the same as in Section 4.

Table D.1 reports the ratio of the long-run trade elasticity, β_{14}^X , to the short-run elasticity, β_0^X , for each process. The main results of Section 4 hold in every column. First, when tariffs are persistent ($\rho = 0.95$ and $\rho = 0.99$), the long-run trade response is smaller under ambiguity than under standard risk, for every degree of ambiguity in both setups. Second, ambiguity makes trade less sensitive to tariff persistence. As ρ rises from 0.8 to 0.99, the ratio rises by 2.3–2.7 under standard risk, depending on the process, but by only 1.5–2.4 under robust control with $\theta = 0.5$ and 0.2–0.9 with $\theta = 0.01$. Third, when tariffs are less persistent ($\rho = 0.8$), strong ambiguity attenuates the long-run trade response. Thus, these results do not depend on the two-state structure of the tariff process.

The magnitudes do depend on the process, in two ways. First, the effects of robust-control ambiguity weaken as the number of states rises. When $\rho = 0.99$ and $\theta = 0.01$, for example, the ratio under robust control is 34% of its standard-risk counterpart in the baseline, but 52% with nine states and a fixed bound, and 49% with nine states and a fixed variance. The max-min results, in contrast, barely change when the bound is fixed,

because the max-min firm's extra weight always falls on the same tariff. My interpretation is that the robust-control firm's worst case depends on the entire distribution of next period's tariff, not just its upper bound. With more states, tariffs move in smaller steps, so a large tariff increase within a single period is less likely, and a given entropy penalty generates less pessimism. Second, the amplification of the long-run response under mild ambiguity when $\rho = 0.8$ is the least robust result. With $\theta = 0.5$, the ratio exceeds its standard-risk counterpart by 0.35 in the baseline, but by only 0.09 with nine states and a fixed bound and 0.03 with nine states and a fixed variance; with $\theta = 0.1$, nine states, and a fixed variance, it falls slightly below. Thus, the conclusions that ambiguity attenuates the response of trade to persistent tariff changes and makes trade less sensitive to tariff persistence are robust to the tariff process, while amplification of the response to transitory tariff changes under mild ambiguity is weaker and does not occur under every alternative tariff process.

Table D.1: Long-run trade elasticities with an AR(1) tariff process

	Baseline	Fixed bound				Fixed variance		
States (N)	2	3	5	9	3	5	9	
Tariff bound (τ_H)	1.30	1.30	1.30	1.30	1.42	1.60	1.85	
Mean tariff	1.15	1.15	1.15	1.15	1.21	1.30	1.42	
S.d. of tariff	0.15	0.11	0.07	0.05	0.15	0.15	0.15	
<i>Panel A: Standard risk and robust control</i>								
$\rho = 0.8$								
Standard risk	1.34	1.44	1.49	1.52	1.62	1.92	2.16	
Robust control, $\theta = 0.5$	1.68	1.66	1.63	1.60	1.86	2.05	2.19	
Robust control, $\theta = 0.1$	1.55	1.62	1.66	1.67	1.75	1.94	2.11	
Robust control, $\theta = 0.05$	1.36	1.50	1.60	1.65	1.54	1.74	1.97	
Robust control, $\theta = 0.01$	1.08	1.15	1.28	1.43	1.12	1.20	1.34	
$\rho = 0.95$								
Standard risk	3.17	3.37	3.55	3.65	3.52	3.90	4.02	
Robust control, $\theta = 0.5$	2.95	3.13	3.34	3.52	3.26	3.52	3.67	
Robust control, $\theta = 0.1$	2.23	2.49	2.77	3.05	2.62	2.96	3.24	
Robust control, $\theta = 0.05$	1.81	2.14	2.46	2.76	2.18	2.57	2.95	
Robust control, $\theta = 0.01$	1.17	1.35	1.63	1.98	1.28	1.45	1.79	
$\rho = 0.99$								
Standard risk	3.63	3.80	4.03	4.20	4.02	4.37	4.49	
Robust control, $\theta = 0.5$	3.17	3.47	3.78	4.04	3.57	3.91	4.16	
Robust control, $\theta = 0.1$	2.39	2.67	3.01	3.40	2.85	3.26	3.61	
Robust control, $\theta = 0.05$	1.99	2.32	2.65	3.01	2.46	2.91	3.33	
Robust control, $\theta = 0.01$	1.23	1.45	1.79	2.17	1.40	1.68	2.20	
<i>Panel B: Max-min</i>								
$\rho = 0.8$								
Max-min, $\alpha = 0.1$	1.38	1.47	1.51	1.53	1.65	1.92	2.12	
Max-min, $\alpha = 0.25$	1.39	1.45	1.47	1.48	1.62	1.85	2.04	
Max-min, $\alpha = 0.5$	1.31	1.33	1.34	1.34	1.48	1.66	1.84	
Max-min, $\alpha = 0.75$	1.17	1.17	1.17	1.17	1.27	1.38	1.52	
$\rho = 0.95$								
Max-min, $\alpha = 0.1$	2.87	3.00	3.07	3.09	3.21	3.50	3.62	
Max-min, $\alpha = 0.25$	2.39	2.46	2.48	2.49	2.73	2.99	3.17	
Max-min, $\alpha = 0.5$	1.79	1.81	1.81	1.81	2.06	2.30	2.55	
Max-min, $\alpha = 0.75$	1.34	1.35	1.35	1.34	1.50	1.67	1.88	
$\rho = 0.99$								
Max-min, $\alpha = 0.1$	3.19	3.29	3.34	3.37	3.54	3.79	3.92	
Max-min, $\alpha = 0.25$	2.55	2.58	2.60	2.60	2.90	3.16	3.38	
Max-min, $\alpha = 0.5$	1.83	1.84	1.84	1.84	2.12	2.38	2.67	
Max-min, $\alpha = 0.75$	1.35	1.35	1.35	1.34	1.52	1.71	1.94	

Notes: Entries are β_{14}^X/β_0^X , the ratio of the long-run to the short-run trade elasticity estimated from (41) on simulated data. Baseline: the two-state process of Section 4 (Figures 2(d) and B.2(d)). Other columns: N -state discretization of a symmetric AR(1) process using the method of Rouwenhorst (1995), with tariffs evenly spaced between $\tau_L = 1$ and τ_H and first-order autocorrelation $2\rho - 1$. Fixed bound: $\tau_H = 1.3$. Fixed variance: $\tau_H = 1 + 0.3\sqrt{N-1}$, which holds the standard deviation of tariffs at its baseline value. Mean and standard deviation of tariffs are computed under the stationary distribution. Other parameters as in Table 2.

E Law of motion for the firm distribution

The distribution of firms over productivity and variable trade costs, $\varphi_t(z, \xi)$, depends on tariffs through the previous period's entry and continuation decisions. Its law of motion is

$$\varphi_t(\mathcal{Z}, \infty) = \sum_{\xi} \left[\int_0^{z^*(\xi, \tau_{t-1})} h(\mathcal{Z}, z) \varphi_{t-1}(z, \xi) dz + \int_0^{\infty} \bar{h}(\mathcal{Z}, z) \varphi_{t-1}(z, \xi) dz \right], \quad (\text{E.1})$$

$$\begin{aligned} \varphi_t(\mathcal{Z}, \xi_H) &= \int_{z^*(\infty, \tau_{t-1})}^{\infty} h(\mathcal{Z}, z) \varphi_{t-1}(z, \infty) dz + \rho_{\xi} \int_{z^*(\xi_H, \tau_{t-1})}^{\infty} h(\mathcal{Z}, z) \varphi_{t-1}(z, \xi_H) dz \\ &\quad + (1 - \rho_{\xi}) \int_{z^*(\xi_L, \tau_{t-1})}^{\infty} h(\mathcal{Z}, z) \varphi_{t-1}(z, \xi_L) dz, \end{aligned} \quad (\text{E.2})$$

$$\varphi_t(\mathcal{Z}, \xi_L) = (1 - \rho_{\xi}) \int_{z^*(\xi_H, \tau_{t-1})}^{\infty} h(\mathcal{Z}, z) \varphi_{t-1}(z, \xi_H) dz + \rho_{\xi} \int_{z^*(\xi_L, \tau_{t-1})}^{\infty} h(\mathcal{Z}, z) \varphi_{t-1}(z, \xi_L) dz, \quad (\text{E.3})$$

where $\mathcal{Z}$ is a typical subset of $\mathbb{R}_{++}$, $h(\mathcal{Z}, z)$ is the probability of surviving and drawing a new productivity in $\mathcal{Z}$ conditional on today's productivity z under the process (34), and $\bar{h}(\mathcal{Z}, z) = [1 - \delta(z)]F_z(\mathcal{Z})$ is the probability of dying and being replaced by a new firm with productivity in $\mathcal{Z}$ conditional on current productivity z . Here F_z is the ergodic productivity distribution from which entrants draw.