Who Can Control Our Firms? National Security and Foreign Ownership Regulation*

Rodrigo Adão
University of Chicago

Olivier Kooi
London School of Economics

March 2026

Abstract

Governments around the world regulate foreign ownership for national security reasons. We develop a theory of foreign ownership regulation based on two hold-up problems. The first is a hold-up problem between firms, which determines ownership structure and firm boundaries. The firm's choice of headquarters location and ownership structure determines its relative use of non-contractible foreign inputs. The second is a geopolitical hold-up problem between governments, which shapes bargaining outcomes in crises. Foreign ownership regulation is optimal because private firms do not internalize that ownership affects reliance on foreign non-contractible inputs and national bargaining power in crises. Regulation is tighter in sectors where output is more valuable in conflict, because defense demand rises or the loss of foreign non-contractible inputs causes larger output declines. Using new data on CFIUS reviews, we show that the United States imposes tighter regulation on sectors with higher defense demand and greater reliance on non-contractible inputs, consistent with the model's predictions.

*Author contacts: rodrigo.adao@chicagobooth.edu, o.b.kooi@lse.ac.uk. We are grateful to Xuhang Ke for outstanding research assistance, to Pol Antràs and John Becko for helpful comments.

1 Introduction

Governments around the world regulate foreign acquisitions of domestic firms. They invoke national security concerns to review private transactions, restrict ownership and control rights, and, in some cases, block acquisitions altogether. Recent examples include the U.S. government's decision to require a golden share as a condition for approving the acquisition of U.S. Steel by Nippon Steel, and the intervention by the Dutch government in the governance of the Chinese-controlled chipmaker Nexperia. More systematically, the Committee on Foreign Investment in the United States (CFIUS) is mandated by Congress to review foreign investments in U.S. businesses with the objective of mitigating risks to national security. The scope of CFIUS regulation has expanded alongside rising geopolitical tensions in recent years, with the number of reviewed transactions in 2022 quadrupling relative to its average level over 2005–2015.

Despite the growing prevalence of such interventions, their economic justification remains unclear. What market failure related to national security leads governments to regulate foreign ownership and control of domestic firms engaged in otherwise voluntary transactions? Why is ownership regulation the preferred instrument, as opposed to alternative policies such as production subsidies or trade taxes? And, given this rationale, how should such regulation be targeted across economic activities?

We address these questions in a framework that combines two hold-up problems arising from contractual frictions. The first is a private-sector hold-up problem between firms, which determines ownership structure and the boundaries of multinational enterprises. The second is a geopolitical hold-up problem between governments, which shapes international bargaining outcomes in times of crisis. While the private-sector hold-up problem determines firms' organizational choices, the geopolitical hold-up problem determines how these choices affect the domestic government's international bargaining position and, in turn, national welfare and optimal policy.

Our central insight is that the optimal design of national-security policy depends on the interaction of contractual frictions behind these two hold-up problems. When production relies on non-contractible inputs, governments cannot directly target firms' reliance on foreign headquarters services using conventional instruments such as production subsidies or trade taxes. Ownership regulation instead emerges as the optimal second-best instrument: by restricting foreign ownership, it reduces firms' reliance on non-contractible foreign headquarter inputs and thereby improves the country's outside option in the event of conflict. This strengthens the government's bargaining position in international negotiations and increases national welfare.

To formalize this insight, we start with a model featuring a private sector hold-up problem that underlies the structure of multinational enterprises. In the spirit of Antràs (2003) and Antràs and Helpman (2004), production in domestic plants requires headquarters services provided by a producer located either domestically or abroad. Firm organization involves two choices: the ownership structure—vertical integration or arm's-length licensing—and the location of headquarters—domestic or foreign. Following Grossman and Hart (1986) and Hart and Moore (1990) (GHM), contractual frictions lead to ex post bargaining under any organizational form, while the ownership structure determines the parties' bargaining power and thereby shapes firms' organizational choices.[1] Vertical integration reallocates bargaining power from plant operators to headquarter producers, increasing the equilibrium provision of headquarters services. Anticipating these outcomes, firms choose their ownership structure and headquarter location to maximize joint private surplus. In equilibrium, all four organizational forms arise due to heterogeneity in the fixed costs of operating under different modes.

In the absence of geopolitical considerations, equilibrium organization choices are constrained efficient, as firms optimally respond to contractual frictions when choosing ownership structure and headquarter location. This establishes a clear benchmark: ownership allocations optimally resolve the private hold-up problem given contractual constraints. Because the provision of headquarters services is non-contractible, the equilibrium allocation reflects the limits imposed by incomplete contracts. These contractual frictions also constrain the government's policy instruments. In particular, the government can condition policy only on verifiable outcomes, such as revenue or organizational form, but not on the non-contractible inputs supplied within firms. Intuitively, if courts cannot verify input provision to enforce private contracts, governments cannot rely on such information to implement economic policy. As a result, the government cannot directly target firms' reliance on headquarters services, whether domestic or foreign.

We then introduce the second hold-up problem, which arises between governments. National governments cannot commit not to resort to coercive actions to obtain concessions from their adversaries. Formally, we consider a geopolitical bargaining game between the domestic and foreign governments, as in Kooi (2024). In the event of a crisis, the foreign government makes a take-it-or-leave-it demand. If the domestic government

[1]This contracting structure follows the property-rights literature. The provision of headquarters services involves non-verifiable, relationship-specific actions, so contracts cannot be conditioned on inputs supplied by each party. This can be justified by the inability of third parties to verify these actions and the possibility of ex post renegotiation (Hart and Moore, 1999; Segal, 1999). Instead, we assume that realized revenue is verifiable and can be contracted upon. As a result, firms optimally choose their scale, but not the ratio of non-contractible inputs.

rejects the offer, negotiations break down, defense spending increases, and international flows of headquarters services are shut down. Because firms differ in their reliance on foreign headquarters services, they experience different output losses in conflict. In equilibrium, conflict does not occur: the foreign government makes an offer that leaves the domestic government indifferent between accepting the demand and entering conflict. Nevertheless, the magnitude of the output loss in conflict determines the domestic government's outside option and thus the equilibrium bargaining outcome.

The interaction between the private-sector and geopolitical hold-up problems overturns the constrained-efficiency benchmark by introducing a wedge between private and social valuations of organizational modes. Firms choose ownership structure and headquarter location to maximize joint private surplus, taking the government's geopolitical bargaining position as given. Therefore they do not internalize how their organizational choices affect the domestic government's bargaining position. From a social perspective, greater reliance on foreign headquarters services weakens the government's bargaining position by increasing the cost of conflict. As a result, organizational modes that rely more heavily on foreign headquarters services—such as vertical integration with foreign headquarters—generate larger geopolitical vulnerabilities and are privately overutilized.

This framework delivers a novel implication for the optimal design of policy in environments with contractual frictions. Because governments face the same contractual limitations as private agents, they cannot directly regulate firms' reliance on non-contractible headquarter inputs. Ownership regulation therefore emerges as the second-best policy: by restricting foreign ownership, the government shifts firms' organizational choices and reduces their reliance on foreign headquarters services, thereby improving the country's outside option in geopolitical bargaining. Stricter regulation applies to organizational forms that rely more heavily on foreign headquarter inputs. Importantly, ownership regulation emerges not because ownership itself generates inefficiencies, but because contractual frictions prevent governments from directly regulating firms' reliance on non-contractible foreign inputs.

Finally, we show how ownership regulation varies with both geopolitical frictions and sectoral characteristics. As geopolitical tensions rise, it is more likely that the government engages in international bargaining, which widens the gap between the private and social valuations of foreign headquarters services. This leads to more severe restrictions on foreign ownership. Across sectors, the intensity of government intervention increases with the marginal utility of consumption in conflict. Restrictions are therefore stricter in sectors that are more important for national defense during conflict and rely more heavily on non-contractible inputs. Interestingly, restrictions are increasing on the interaction of these two determinants, as the impact of output losses on the marginal utility of consumption increases with the level of defense spending during conflict.

While our model is not the only possible national security rationale for foreign ownership regulations, we view its logic as quite general. Ownership structure is not directly payoff-relevant in geopolitical interactions. Instead, it matters because it shapes private decision-making, which in turn affects a government's bargaining position in international negotiations. To illustrate the broader applicability of this mechanism, we study an extension based on an agency theory of firm ownership. In this environment, the government values domestic ownership because it facilitates the punishment of domestic firms that take non-contractible actions weakening the government's bargaining position. Ownership again emerges as a second-best instrument for influencing private decisions that have national security consequences.

In the final part of the paper, we confront the model's predictions with data on how the United States regulates foreign ownership and control of domestic firms. We construct a new dataset based on cases reviewed by CFIUS that measures the intensity of foreign ownership restrictions across sectors. Our empirical strategy exploits cross-sector variation in regulatory intensity and relates it to observable proxies for the theoretical determinants of optimal intervention. Consistent with the model's predictions, regulatory intensity is higher in sectors with greater military consumption and greater reliance on non-contractible inputs, as measured by Nunn (2007), and is especially high in sectors where these two characteristics coincide. Together, these two measures explain 53 percent of the variation of our measure of regulatory intensity across four-digit manufacturing industries in the United States.

Related Literature. The interaction between private contractual frictions and geopolitical bargaining determines the set of feasible policy instruments for governments and, consequently, the optimal design of national-security policy. On the one hand, the "macro" hold-up problem between governments reflects the absence of enforceable agreements between sovereign states and implies that economic dependencies can affect international bargaining outcomes (Waltz, 1959; McLaren, 1997). On the other hand, the "micro" hold-up problem between private agents provides the foundation for the property-rights theory of the firm, in which ownership structure determines bargaining power and investment incentives under incomplete contracts (Grossman and Hart, 1986; Hart and Moore, 1990). Our contribution is to study the interaction between these two hold-up problems. While ownership structure is chosen to address private hold-up concerns, it also determines firms' reliance on foreign inputs and thereby affects the domestic government's outside option in geopolitical bargaining. As a result, firm organization choices that are privately optimal need not be socially optimal once geopolitical hold-up considerations are taken into account.

Our analysis relates to a growing geoeconomics literature that studies national security motives for economic policy in environments where countries can be coerced through trade, financial, or production linkages (e.g., Clayton et al., 2023, 2024; Becko and O'Connor, 2024; Kooi, 2024; see Mohr and Trebesch, 2025 for a review). This literature shows that governments intervene to reduce strategic vulnerabilities by aligning private and social returns to production across sectors. When the relevant determinants of geopolitical outcomes are directly contractible, this logic yields a Pigouvian prescription: production is encouraged in sectors that are more valuable in conflict and discouraged in sectors that expose the country to geopolitical risk.

Theoretically, we show that this logic breaks down when contractual frictions within firms shape the relevant determinants of geopolitical outcomes. The key margin is firms' reliance on non-contractible foreign inputs, which affects the country's outside option in geopolitical bargaining but is not directly observable or contractible. As a result, governments cannot directly target the underlying source of vulnerability. Ownership regulation then emerges as the optimal second-best instrument: it shifts firms' organizational form, which determines their exposure to foreign inputs and, in doing so, indirectly affects the country's exposure to geopolitical risk. Empirically, we provide evidence that foreign ownership regulation in the United States is consistent with the model's predictions that regulatory intensity increases with defense demand and non-contractible input intensity across sectors. This approach complements the existing literature by confronting the theory with observed policy instruments and by explaining empirical patterns that models without a theory of the firm cannot account for.[2]

Our paper places the private-sector hold-up problem at center stage to generate endogenous organizational choices. We build on models of the boundaries of multinational firms, such as Antràs (2003) and Antràs and Helpman (2004), and adopt a contracting structure that delivers constrained efficiency in the absence of geopolitical hold-up considerations. Within this framework, foreign ownership leads to deeper integration but also increases the cost of losing access to foreign headquarters services. This prediction is consistent with evidence on the low substitutability of headquarter inputs during large production disruptions abroad (Boehm et al., 2019). More broadly, our analysis shows

[2]Most of the literature studies optimal policy in isolation from military considerations. Exceptions include Kooi (2024) and Alekseev and Lin (2026). Incorporating military considerations into the model allows us to account for the evidence that it is a central determinant of foreign ownership regulation.

how the non-contractibility of headquarter inputs links firms' organizational choices to geopolitical bargaining outcomes and makes ownership regulation an effective second-best policy instrument.[3]

2 Optimal International Firm Boundaries With Private-Sector Hold-up

In this section, we introduce a model in which the structure of firm ownership arises from a hold-up problem between domestic suppliers of specialized inputs and providers of headquarters services located either domestically or abroad, as in Antràs and Helpman (2004). We show that the government cannot improve on the equilibrium structure of firm ownership whenever it is subject to the same contractual frictions faced by private agents in choosing firm boundaries (Grossman and Hart, 1986; Hart and Moore, 1990).

2.1 A Model of International Firm Boundaries

Preferences. We consider a single country with a representative household that inelastically supplies L labor units, and derives utility from consuming homogeneous goods, q_g for $g \in \{0, 1, \ldots, G\}$:

$$U = q_0 + \sum_{g=1}^{G} \frac{\sigma_g}{\sigma_g - 1} (q_g)^{\frac{\sigma_g - 1}{\sigma_g}}, \tag{1}$$

where $\sigma_g > 0$ is the elasticity of demand for good g.

Home technology of consumption goods. The economy produces the consumption goods $g \in \{0, 1, \ldots, G\}$. Good 0 is freely traded internationally and produced competitively using only domestic labor under a linear technology, $y_0 = l_0$. Free trade in good 0 pins down the domestic wage and makes it the numéraire.

The remaining goods, $g \in \{1, \ldots, G\}$, are non-traded and produced domestically. For each good g, the representative household owns a unit mass of competitive firms indexed by i. All firms producing good g have the same production technology,

$$y_g(i) = \left[(l_g(i))^{1-\alpha_g} \left((h_g(i))^{\eta_g^h} (m_g(i))^{\eta_g^m} \right)^{\alpha_g} \right]^{\psi_g}, \tag{2}$$

[3] Our framework also relates to a broader literature studying the interaction between multinational firms and international political outcomes (e.g., Gilpin, 1975; Wellhausen, 2015). We provide a microfounded mechanism through which firm ownership affects national bargaining power and derive the optimal policy response when these effects arise from non-contractible production decisions.

where $l_g(i)$ denotes contractible labor, $m_g(i)$ denotes non-contractible intermediate inputs supplied by specialized domestic operators, and $h_g(i)$ denotes non-contractible headquarters services.

This production structure has three key features. First, $\psi_g \in (0,1)$ governs decreasing returns to scale at the firm level, ensuring that each good is produced by multiple firms under perfect competition and preserving a meaningful role for firm-level organizational choices. Second, $\alpha_g \in (0,1)$ determines the importance of non-contractible inputs in production, thereby controlling the extent to which contractual frictions affect firm behavior. Third, the composite non-contractible input combines intermediate inputs supplied by specialized domestic operators, $m_g(i)$, and headquarters services, $h_g(i)$. The parameter $\eta_g^h \in (0,1)$ measures the intensity of headquarters services, with $\eta_g^m + \eta_g^h = 1$. Headquarters services are provided by either a domestic or foreign specialized producer, indexed by $\ell \in \{D, F\}$, and are perfect substitutes in production.

Headquarters services and domestic inputs. Specialized operators produce intermediate inputs using domestic labor under a linear technology, with a marginal cost of c_g^m. Intermediate inputs are not traded internationally and must be sourced domestically.

Headquarters services are produced by specialized headquarter producers for each good g. Domestic headquarter producers require $c_{g,D}^h$ units of domestic labor to produce one unit of headquarters services, implying a marginal cost of $c_{g,D}^h$. These services must be sold domestically. Alternatively, firms may source headquarters services from foreign producers, who can supply these services at a marginal cost of $c_{g,F}^h$.

Organizational modes. In the tradition of the GHM theory of the firm, contractual frictions give rise to a private-sector hold-up problem. To structure its relationship with headquarters providers and specialized input suppliers, each firm i chooses an organizational mode that entails two decisions. The first is whether to rely on domestic ($\ell = D$) or foreign ($\ell = F$) providers of headquarters services. The second is whether to vertically integrate with headquarters producers ($o = V$) or engage in arm's-length outsourcing ($o = A$). Each organizational mode entails a fixed cost in terms of domestic labor, denoted by $F_{g,o\ell}(i)$. The negative of this cost i draws from a Type I Extreme Value (Gumbel) distribution with location parameter $\bar{F}_{g,o\ell}$ and dispersion parameter $1/\nu_g$. Given its draw of fixed costs, firm i chooses its organizational mode to maximize joint surplus net of fixed costs, taking into account how ownership structure affects input provision.[4]

[4]As in Antràs and Helpman (2004), the fixed cost captures the fact that different organizational modes may require distinct upfront setup costs. We depart from Antràs and Helpman (2004) by assuming that

The hold-up problem arises because headquarters services and intermediate inputs are non-verifiable and relationship-specific. Contracts cannot condition on the realized provision of $(m_g(i), h_g(i))$. As a result, the parties cannot commit ex ante to the input composition and instead go through an ex post bargaining stage (Hart and Moore, 1999; Segal, 1999).

We assume the following contracting structure. In line with Antràs and Helpman (2004), before investment, the firm offers lump-sum transfers $S_g^h(i)$ and $S_g^m(i)$ to the headquarters provider and the specialized domestic operators, respectively. These transfers ensure participation, and we normalize the outside option of each supplier to zero. Departing from Antràs and Helpman (2004), we assume a richer contract in which the firm can also commit to a joint revenue-based payment $s_g(i)R_g(i)$ to the non-contractible input providers, contingent on verifiable revenue. However, it cannot condition payments on the composition of inputs provided by each party, reflecting their non-verifiable nature. Note that, in Antràs and Helpman (2004), the firm is restricted to select $s_g(i) = 1$.

After production decisions are made, the parties engage in Nash bargaining over the division of this joint payment. Because the revenue-based payment scales proportionally with output, it affects incentives to supply non-contractible inputs equally for both parties without altering their relative bargaining positions. Intuitively, the firm can deposit $S_g(i)$ in an escrow account and use these funds to finance a revenue-proportional subsidy, such that $S_g(i) \geq s_g(i)R_g(i)$. Funds remaining in escrow are wasted from the firm's perspective.

Ownership structure determines the allocation of bargaining power in ex post negotiations. Under organizational mode $o \in \{V, A\}$, the headquarters producer and the specialized domestic operators receive shares $\beta_{g,o}^h$ and $\beta_{g,o}^m$ of the joint payment to non-contractible inputs, with $\beta_{g,o}^h + \beta_{g,o}^m = 1$. Following Antràs (2003) and Antràs and Helpman (2004), these shares are exogenous parameters that reflect the allocation of residual control rights implied by the ownership structure. In particular, vertical integration is associated with a higher share received by headquarters producers, $\beta_{g,V}^h > \beta_{g,A}^h$.[5]

Government. The government designs policy under the same contractual frictions faced by private agents. Accordingly, the government can set regulations that depend on the firm's organizational mode and contractible outcomes, but cannot condition regulation

firms are heterogeneous only with respect to these fixed costs. While generating heterogeneity in organizational choices within a sector, this ensures that aggregate differences arise from organizational choices themselves, rather than from underlying firm characteristics.

[5]Our bargaining protocol follows from the microfoundation developed in Antràs and Helpman (2004). In this microfoundation, vertical integration increases the fraction of output the headquarters can appropriate when the two parties do not agree, thereby strengthening the headquarters' ex-post bargaining power.

on the firm's composition of non-contractible inputs. We model regulation as revenue-neutral wedges, which allow the government to directly control firm's perceived variables, without generating government revenue.[6] Specifically, given the firm's organizational mode, the government sets a wedge on the firm's revenue $\tau^R_{g,o\ell}$, its payments to contractible inputs $\tau^L_{g,o\ell}$, as well as wedges on location and ownership choices $\tau^O_{g,o\ell}$. The government uses these organizational wedges to directly affect foreign ownership and, thus, replicate foreign ownership regulation. In addition, the government has an exogenous demand for good g of d_g, which is used for national defense. The government budget is balanced, so that any net tax revenue is rebated to the representative household through a lump-sum transfer, T.

2.2 Competitive equilibrium

We now characterize the competitive equilibrium in this economy.

Non-contractible input choices. Firms are heterogeneous only with respect to organizational fixed costs. Therefore, all firms that select the same organizational mode make identical input choices. To simplify notation, we drop the firm index i and index variables by organizational mode (o, ℓ).

Given an organizational mode, each non-contractible input supplier chooses its input taking as given: (i) the other supplier's input choice; (ii) the firm's choice of contractible labor, $l_{g,o\ell}$; (iii) the firm's internal revenue-sharing parameter, $s_{g,o\ell}$; (iv) the government's wedges, $\{\tau^R_{g,o\ell}, \tau^L_{g,o\ell}, \tau^O_{g,o\ell}\}$; and (v) the domestic price of good g, p_g. Thus, $h_{g,o\ell}$ and $m_{g,o\ell}$ solve:

$$h_{g,o\ell} = \arg\max_h \{\beta^h_{g,o} s_{g,o\ell} p_g (1 - \tau^R_{g,o\ell}) \left(l_{g,o\ell}\right)^{(1-\alpha_g)\psi_g} \left((h)^{\eta^h_g} (m_{g,o\ell})^{\eta^m_g}\right)^{\alpha_g \psi_g} - c^h_{g,\ell} h\}$$
$$m_{g,o\ell} = \arg\max_m \{\beta^m_{g,o} s_{g,o\ell} p_g (1 - \tau^R_{g,o\ell}) \left(l_{g,o\ell}\right)^{(1-\alpha_g)\psi_g} \left((h_{g,o\ell})^{\eta^h_g} (m)^{\eta^m_g}\right)^{\alpha_g \psi_g} - c^m_g m\}.$$

[6]It is well known that many different policy instruments can decentralize a given allocation. Implementing allocations through wedges rather than taxes is common in the macroprudential policy literature (see e.g Costinot et al. (2014) and Farhi and Werning (2016)). We prefer regulatory wedges over taxes for two reasons. First, regulatory wedges have a direct mapping to the economic incentives faced by firms, which simplifies exposition and comparative statics. Second, in practice, most governments regulate foreign ownership, but do not set taxes conditional on foreign ownership. Section 3.4 shows that our main insights hold if we allow the government to choose taxes instead of revenue-neutral regulatory wedges.

The solution of this system is

$$h_{g,o\ell} = \beta^h_{g,o}(\eta^h_g/c^h_{g,\ell}) \left(\kappa_{g,o\ell} s_{g,o\ell} p_g (1-\tau^R_{g,o\ell})(l_{g,o\ell})^{(1-\alpha_g)\psi_g}\right)^{\frac{1}{1-\alpha_g\psi_g}} \tag{3}$$

$$m_{g,o\ell} = \beta^m_{g,o}(\eta^m_g/c^m_g) \left(\kappa_{g,o\ell} s_{g,o\ell} p_g (1-\tau^R_{g,o\ell})(l_{g,o\ell})^{(1-\alpha_g)\psi_g}\right)^{\frac{1}{1-\alpha_g\psi_g}} \tag{4}$$

where $\kappa_{g,o\ell} \equiv \alpha_g\psi_g(\beta^h_{g,o}\eta^h_g/c^h_{g,\ell})^{\eta^h_g\alpha_g\psi_g}(\beta^m_{g,o}\eta^m_g/c^m_g)^{\eta^m_g\alpha_g\psi_g}$.

Expressions (3)–(4) highlight two features. First, ex post bargaining distorts the firm's relative composition of non-contractible input. Under full contractibility, both inputs would receive their marginal revenue products, and the input ratio would depend only on technological intensity and relative marginal costs. This efficient ratio is attained under incomplete contracting if and only if $\beta^h_{g,o} = \beta^m_{g,o}$. When $\beta^h_{g,o}$ is higher—as under vertical integration—firms rely more heavily on headquarters services. Second, the revenue-based payment parameter $s_{g,o\ell}$ scales both inputs proportionally and therefore affects production scale, but not their relative composition.

Firm's optimal internal contract. Given its organizational mode, the firm chooses its internal contract to maximize total surplus. Because revenue-based payments must be self-funded, the firm sets the minimum necessary deposit, $S_{g,o\ell} = s_{g,o\ell}(1-\tau^R_{g,o\ell})p_g y_{g,o\ell}$. To ensure participation of non-contractible suppliers, the firm makes lump-sum transfers such that participation constraints bind: $S^h_{g,o\ell} = c^h_{g,\ell}h_{g,o\ell} - \beta^h_{g,o}S_{g,o\ell}$ and $S^m_{g,o\ell} = c^m_g m_{g,o\ell} - \beta^m_{g,o}S_{g,o\ell}$. The firm therefore chooses $(s_{g,o\ell}, l_{g,o\ell})$ to maximize total surplus accounting for ex post non-contractible input choices:

$$(s_{g,o\ell}, l_{g,o\ell}) \in \arg\max_{\{s,l\}} (1-\tau^R_{g,o\ell})p_g y_{g,o\ell} - (1+\tau^L_{g,o\ell})l - c^h_{g,\ell}h_{g,o\ell} - c^m_g m_{g,o\ell}$$

$$\text{subject to} \quad (y_{g,o\ell}, h_{g,o\ell}, m_{g,o\ell}) \text{ given by (2)–(4).}$$

It is optimal for the firm to set the revenue sensitivity of payments to non-contractible input suppliers above one, $s_{g,o\ell} = (\eta^h_g\beta^h_{g,o} + \eta^m_g\beta^m_{g,o})^{-1} > 1$. Intuitively, although the firm cannot correct the distortion in the relative composition of non-contractible inputs, it uses the revenue-based payment to expand production scale and mitigate the standard underinvestment distortion generated by ex post bargaining. Even when $\beta^h_{g,o} = \beta^m_{g,o}$, so that the input mix is efficient, the firm sets $s_{g,o\ell} = 2$ to eliminate the scale distortion.

The firm's output with regulatory wedges is given by

$$y_{g,o\ell} = \left(\rho_g \frac{(1-\tau^R_{g,o\ell})p_g}{(1+\tau^L_{g,o\ell})^{1-\alpha_g}} \left(\frac{(\beta^h_{g,o}\eta^h_g / c^h_{g,\ell})^{\eta^h_g} (\beta^m_{g,o}\eta^m_g / c^m_g)^{\eta^m_g}}{\beta^h_{g,o}\eta^h_g + \beta^m_{g,o}\eta^m_g} \right)^{\alpha_g} \right)^{\frac{\psi_g}{1-\psi_g}} \tag{5}$$

with $\rho_g \equiv \psi_g \alpha_g^{\alpha_g} (1-\alpha_g)^{1-\alpha_g}$.

Organizational mode. The firm chooses its organizational mode to maximize profits net of fixed costs and regulatory wedges. Given the distribution of fixed costs, the share of firms producing good g that select organizational mode (o, ℓ) is

$$P_{g,o\ell} = \frac{\exp\left(\nu_g(\pi_{g,o\ell} + \bar{F}_{g,o\ell} - \tau^O_{g,o\ell})\right)}{\sum_{o',\ell'} \exp\left(\nu_g(\pi_{g,o'\ell'} + \bar{F}_{g,o'\ell'} - \tau^O_{g,o'\ell'})\right)} \tag{6}$$

with realized profits $\pi_{g,o\ell} = p_g y_{g,o\ell} - c^h_{g,\ell} h_{g,o\ell} - c^m_g m_{g,o\ell} - l_{g,o\ell}$. Note that regulatory wedges only enter profits indirectly through their influence on firms' decisions. This follows from the assumption that the wedges affect directly the firm's perceived revenue and wage, but do not generate government revenue.[7]

Final demand. The utility maximization problem of the representative household implies that, for each non-traded good $g \in \{1, ..., G\}$, the final demand is given by

$$q_g = (p_g)^{-\sigma_g}. \tag{7}$$

For good 0, budget balance implies that $q_0 = E - \sum_{g=1}^G (p_g)^{1-\sigma_g}$, with E the aggregate domestic expenditure.

Market clearing. For each non-traded good $g \in \{1, ..., G\}$, the domestic price p_g is determined by domestic market clearing:

$$q_g + d_g = y_g \equiv \sum_{o,\ell} P_{g,o\ell} y_{g,o\ell}, \tag{8}$$

where y_g is the domestic output of g, with $y_{g,o\ell}$ given by (5), $P_{g,o\ell}$ by (6), and q_g by (7).

[7]Intuitively, our wedges replicate an organization-specific bundle of taxes on contractible revenue and labor and lump-sum transfers that exactly offset each other. As such, our wedges affect the firm's marginal product of inputs, but do not affect its overall profits. Section 3.4 formalizes this equivalence between wedges and taxes.

We assume that labor supply is sufficiently large so that the domestic country produces a strictly positive amount of good 0. The labor market clearing then determines its equilibrium output, y_0, and the budget constraint its equilibrium consumption, q_0.[8]

2.3 Constrained Efficiency of International Firm Boundaries

This section establishes that the competitive equilibrium is constrained efficient. That is, given the same contractual frictions faced by private agents, the government cannot raise aggregate welfare using policy instruments that depend only on contractible outcomes. Using the numeraire technology and market clearing, the representative household's utility can be written as

$$U = L + \sum_{g=1}^{G} \frac{\sigma_g}{\sigma_g - 1}(y_g - d_g)^{\frac{\sigma_g - 1}{\sigma_g}} - \sum_{g,o,\ell} P_{g,o\ell}(l_{g,o\ell} + c_g^m m_{g,o\ell} + c_{g,\ell}^h h_{g,o\ell} + \nu_g^{-1} \log P_{g,o\ell}) \quad (9)$$

where the last term is the average fixed cost of firms selecting organizational mode (o, ℓ).

The government chooses allocations subject to the same contractual constraints that determine firms' input choices. In particular, because non-contractible inputs are chosen under ex post bargaining, feasible allocations must satisfy the equilibrium input ratio implied by (3)–(4). The government's problem is therefore

$$\max_{\{P_{g,o\ell}, y_{g,o\ell}, l_{g,o\ell}, h_{g,o\ell}, m_{g,o\ell}\}} U \quad \text{subject to}$$

$$\sum_{o,\ell} P_{g,o\ell} = 1, \quad h_{g,o\ell} = \frac{\beta_{g,o}^h \eta_g^h c_g^m}{\beta_{g,o}^m \eta_g^m c_{g,\ell}^h} m_{g,o\ell}, \quad y_{g,o\ell} = \left[(l_{g,o\ell})^{1-\alpha_g} \left((h_{g,o\ell})^{\eta_g^h} (m_{g,o\ell})^{\eta_g^m} \right)^{\alpha_g} \right]^{\psi_g}$$

Appendix A shows that the solution to this problem coincides with the competitive equilibrium when $\tau_{g,o\ell}^R = \tau_{g,o\ell}^L = \tau_{g,o\ell}^O = 0$, which establishes the following result.

Proposition 1. *The competitive equilibrium is constrained efficient. In the benchmark without geopolitical considerations, it is optimal for the government to set* $\tau_{g,o\ell}^R = \tau_{g,o\ell}^L = \tau_{g,o\ell}^O = 0$ *for all* (g, o, ℓ).

The proposition implies that there is no role for policy in this benchmark economy without geopolitical considerations. When the government and firms face the same contractual frictions, private agents choose organizational modes efficiently given contractual constraints. Intuitively, because policy cannot be made contingent on the compo-

[8] Note that trade balance in equilibrium implies that the domestic country's exports of the numeraire good are equal to its payments to foreign providers of headquarters services, $y_0 - q_0 = \sum_{g,o} P_{g,oF} c_{g,F}^h h_{g,oF}$.

sition of non-contractible inputs, the government cannot directly undo the within-firm distortion in relative input use induced by ex post bargaining.

Alternative contractual structure. Constrained efficiency depends crucially on the assumption that both firms and the government can write contracts contingent on revenue. If firms cannot condition payments on revenue (i.e., $s_{g,o\ell} = 1$ as in Antràs and Helpman (2004)) while the government can, then it is optimal to use revenue wedges to correct the standard underinvestment distortion implied by ex post bargaining: $\tau^R_{g,o} = 1 - (\eta^h_g \beta^h_{g,o} + \eta^m_g \beta^m_{g,o})^{-1}$ for $\ell \in \{D, F\}$. Note that policy does not depend on the location of headquarters, since it does not affect the degree of underinvestment implied by ex post bargaining within the firm.

3 Optimal International Firm Boundaries With Geopolitical Hold-up

In this section, we extend the model to incorporate a geopolitical hold-up problem in which the Domestic country can be coerced through the threat of conflict, as in Kooi (2024). This generates a motive for the Domestic government to intervene in international firm boundaries, even when it faces the same contractual frictions as private agents. As in Grossman and Hart (1986) and Hart and Moore (1990), the source of the hold-up problem is incomplete contracting. Here, however, the friction reflects the absence of an authority above the nation-state capable of enforcing international agreements, rather than non-verifiable contingencies.

3.1 A Model of International Firm Boundaries With Geopolitical Hold-up Between National Governments

Geopolitical bargaining game. We consider a three-stage game between the governments of Domestic and Foreign.

Stage 1. Domestic firms choose their organizational modes and input allocations, anticipating the possibility of future geopolitical conflict but ignoring the impact of their choices on geopolitical outcomes. They also make transfers to input suppliers to satisfy participation constraints.

Stage 2. With probability $1-\theta$, the economy remains in the geopolitical outcome of Peace ($Z = P$), as described in Section 2. With probability θ, a geopolitical crisis occurs in which the Foreign government attempts to coerce the Domestic government. Coercion takes the form of an ultimatum demand for a transfer $\mathcal{T}$ of the traded good 0. If the Domestic government accepts, it transfers $\mathcal{T}$ units of good 0 to the Foreign country and the outcome remains Peace ($Z = P$). If it rejects, the two countries enter the geopolitical outcome of Conflict ($Z = C$).

Conflict differs from Peace in two ways. First, the government increases its demand for goods used in national defense, with fixed demand $d_g^C > d_g$. This captures the strategic relevance of goods during conflict. Second, domestic firms lose access to foreign headquarters services. However, firms have access to a backup technology that allows them to convert final output into headquarters services given their choices of inputs and organizational mode in Stage 1, $(m_{g,oF}, l_{g,oF})$. This generates heterogeneity in firms' exposure to geopolitical disruptions depending on their reliance on foreign headquarters services.

Stage 3. Production and consumption are realized. If an ultimatum is made and accepted, the transfer $\mathcal{T}$ is paid.

Conflict output. Firms that source headquarters services domestically are unaffected by conflict and continue to produce as described in Section 2. Firms that rely on foreign headquarters services must instead produce headquarters services internally by diverting resources from production. Given their choices in Stage 1 ($m_{g,oF}, l_{g,oF}$), firms use a backup technology to generate headquarters services h, yielding net output of

$$y_{g,oF}(h) = \left(l_{g,oF}^{1-\alpha_g}\left(h^{\eta_g^h}(m_{g,oF})^{\eta_g^m}\right)^{\alpha_g}\right)^{\psi_g} - c_g^C h,$$

where c_g^C is a conversion cost parameter capturing the inefficiency of producing headquarters services internally during conflict.

Firms then choose h to maximize net output in conflict. The solution to this problem implies output in Conflict as a function of $(m_{g,oF}, l_{g,oF})$:

$$y_{g,oF}^C = \rho_g^C\left((m_{g,oF})^{\eta_g^m\alpha_g}(l_{g,oF})^{1-\alpha_g}\right)^{\frac{\psi_g}{1-\eta_g^h\alpha_g\psi_g}} \tag{10}$$

where $\rho_g^C \equiv (1-\eta_g^h\alpha_g\psi_g)(\eta_g^h\alpha_g\psi_g/c_g^C)^{\eta_g^h\alpha_g\psi_g/(1-\eta_g^h\alpha_g\psi_g)}$.

Two features of the backup technology are relevant for our purposes. First, when c_g^C is sufficiently high, output in Conflict is lower than in Peace. We impose this condition so that firms relying on foreign headquarters services experience production losses during conflict. Second, output in Conflict is increasing in both $m_{g,oF}$ and $l_{g,oF}$. Consequently, among producers of good g that rely on foreign headquarters services, vertically integrated firms experience larger output losses during conflict because they depend more heavily on headquarters inputs.

Geopolitical outcome. In equilibrium of the geopolitical ultimatum game, the Foreign government demands the largest transfer that the Domestic government is willing to accept in order to remain in Peace. Such a transfer leaves the Domestic government indifferent between Peace and Conflict. The Domestic government accepts the offer, and the equilibrium geopolitical outcome is Peace

Let y_g^C denote domestic output of good g in Conflict, while y_g denotes output in Peace. The equilibrium transfer is therefore

$$\mathcal{T} = \sum_{g=1}^{G} \frac{\sigma_g}{\sigma_g - 1} \left((y_g - d_g)^{\frac{\sigma_g - 1}{\sigma_g}} - (y_g^C - d_g^C)^{\frac{\sigma_g - 1}{\sigma_g}} \right). \tag{11}$$

The transfer captures the change in utility generated by conflict. This loss arises both from production disruptions caused by the loss of access to foreign headquarters services (i.e., $y_{g,oF}^C < y_{g,oF}$) and from the increase in defense requirements (i.e., $d_g^C > d_g$). Note that consumption of good 0 is only affected by the geopolitical outcome through the transfer, since input choices are predetermined in Stage 1 and payments to foreign headquarters providers are financed through escrow accounts.[9]

National security concerns. The ultimatum game gives rise to national security concerns about over-reliance on non-contractible foreign headquarters services. When selecting their organizational modes and internal payment structures, firms take as given the geopolitical outcome of Peace, the international transfer $\mathcal{T}$, and the regulatory wedges $\{\tau_{g,o\ell}^R, \tau_{g,o\ell}^L, \tau_{g,o\ell}^O\}$. However, these choices affect consumption in the event of conflict and therefore influence geopolitical bargaining power and the equilibrium international transfer in Peace, as shown in (11). Since the transfer must ultimately be financed through taxation, these private choices generate an aggregate fiscal externality that firms do not internalize. Consequently, resilience to conflict has value even though conflict never materializes, because the possibility of conflict shapes international bargaining outcomes.

[9] Alternatively, under the assumption that foreigners do not receive payments in the event of conflict, the transfer must be adjusted by the net exports of good 0 associated with payments for headquarters services during Peace, as defined in footnote 8.

The government therefore designs policy to reduce the equilibrium transfer by lowering exposure to conflict.

3.2 International Firm Boundaries With National Security Concerns

We return to the government's problem of designing economic policy, now accounting for the outcome of geopolitical negotiations. The government considers the expected utility cost of the international transfer, $\theta\mathcal{T}$, with $\mathcal{T}$ given by (11), as well as the output loss during conflict, determined by the conflict technology in (10). The government's problem is therefore

$$\max_{\{P_{g,o\ell},y_{g,o\ell},l_{g,o\ell},h_{g,o\ell},m_{g,o\ell}\}} U+\theta\sum_{g=1}^{G}\frac{\sigma_g}{\sigma_g-1}\left(\left(y_g^C-d_g^C\right)^{\frac{\sigma_g-1}{\sigma_g}}-\left(y_g-d_g\right)^{\frac{\sigma_g-1}{\sigma_g}}\right)\quad\text{subject to}$$

$$\sum_{o,\ell}P_{g,o\ell}=1,\quad h_{g,o\ell}=\frac{\beta_{g,o}^{h}\eta_g^h c_g^m}{\beta_{g,o}^{m}\eta_g^m c_{g,\ell}^h}m_{g,o\ell},\quad y_{g,o\ell}=\left[\left(l_{g,o\ell}\right)^{1-\alpha_g}\left(\left(h_{g,o\ell}\right)^{\eta_g^h}\left(m_{g,o\ell}\right)^{\eta_g^m}\right)^{\alpha_g}\right]^{\psi_g},$$

$$y_{g,oD}^C=y_{g,oD},\quad y_{g,oF}^C=\rho_g^C\left(\left(m_{g,oF}\right)^{\eta_g^m\alpha_g}\left(l_{g,oF}\right)^{1-\alpha_g}\right)^{\frac{\psi_g}{1-\eta_g^h\alpha_g\psi_g}}.$$

The government faces both "micro" and "macro" contractual frictions. As in Section 2, it must choose a non-contractible input composition that satisfies the ex post bargaining problem between suppliers within firms, as captured by the constraints in the second row. However, with probability θ, it must also pay the international transfer resulting from ex post bargaining between governments, which depends on the conflict technology captured by the constraints in the last row.

The key source of market failure is that private agents make decisions taking the international transfer during geopolitical crises as given. However, these decisions affect the magnitude of the transfer through the reduction in consumption in the event of conflict. For each good g, this effect is captured by the change in the marginal value of consumption, reflected in the price change in the off-equilibrium conflict outcome,

$$\lambda_g\equiv\frac{p_g^C}{p_g}=\left(\frac{y_g^C-d_g^C}{y_g-d_g}\right)^{-1/\sigma_g}. \tag{12}$$

Since domestic firms' output does not change during conflict, the government only needs to use its policy instruments to induce these firms to internalize the reduction in consumption during conflict. In particular, it sets revenue wedges to increase production in goods that become more valuable in conflict, as stated in the following proposition.

Proposition 2. *For firms with domestic headquarters, the optimal wedges are given by*

$$\tau^R_{g,oD} = \bar{\tau}^O_{g,oD} = \theta\left(1 - \lambda_g\right) \leq 0 \quad \textit{and} \quad \tau^L_{g,oD} = 0, \tag{13}$$

where $\bar{\tau}^O_{g,o\ell} \equiv \tau^O_{g,o\ell} / p_g y_{g,o\ell}$.

Two comments are in order. First, the government encourages production equally across all firms with domestic headquarters, so it sets identical wedges for firms that are vertically integrated and for those that are not. Second, there is no motive to distort input use for firms with domestic headquarters, since they only need to internalize that their production scale is suboptimal in the off-equilibrium conflict state.

Proposition 2 connects naturally to the broader geoeconomics literature on national security and production distortions (e.g., Clayton et al., 2023, 2024; Becko and O'Connor, 2024; Kooi, 2024; see Mohr and Trebesch, 2025 for a review). In these environments, the relevant margins are contractible, so governments can directly intervene to align private and social returns across goods. The optimal policy therefore follows a Pigouvian logic: production is encouraged in sectors that become more valuable in the conflict state, as captured here by λ_g. Proposition 2 shows that, in our setting, this logic can be implemented through wedges on contractible outcomes, which raise production in high-λ_g goods.[10] This allocation is first-best in an environment without foreign headquarters, where geopolitical risk operates only through changes in the marginal value of consumption and can be fully addressed through interventions on contractible variables.

This logic breaks down once firms rely on foreign non-contractible inputs that affect conflict outcomes. In this case, the government must account for the consumption loss during conflict induced by the loss of access to foreign headquarters services,

$$\Delta_{g,oF} \equiv \frac{y^C_{g,oF}}{y_{g,oF}}. \tag{14}$$

Ownership structure determines the magnitude of this output loss. Due to ex post bargaining within the firm, vertically integrated firms rely more intensively on foreign headquarters services, implying that $\Delta_{g,VF} < \Delta_{g,AF} < 1$.

The following proposition characterizes the optimal wedges for firms with foreign headquarters. These wedges induce firms to internalize how their organizational choices

[10] In our framework, this takes the form of wedges on revenue rather than directly on inputs. Because revenue scales proportionally with all inputs, these wedges affect production scale without distorting the composition of contractible and non-contractible inputs. This is equivalent to input subsidies in environments without non-contractible inputs, for example as in Kooi (2024).

affect the marginal value of consumption in the conflict state and, therefore, the international transfer in a geopolitical crisis. The government uses three instruments: a revenue wedge $\tau^R_{g,oF}$ that adjusts the scale of firms across organizational modes, a labor wedge $\tau^L_{g,oF}$ that distorts input composition toward contractible labor, and an organizational wedge $\tau^O_{g,oF}$ that implements the optimal share of firms operating under each mode.

Proposition 3. *For firms with foreign headquarters, the optimal wedges are given by*

$$\tau^R_{g,oF} = \theta\left(1 - \lambda_g \Delta_{g,oF} \frac{\eta^m_g}{1 - \eta^h_g \psi_g \alpha_g}\right), \tag{15}$$

$$\bar{\tau}^O_{g,oF} = \theta\left(1 - \lambda_g \Delta_{g,oF}\right), \tag{16}$$

$$\frac{\tau^L_{g,oF}}{1 + \tau^L_{g,oF}} = -\theta \frac{\lambda_g \Delta_{g,oF}}{1 - \tau^R_{g,oF}} \frac{\eta^h_g}{1 - \eta^h_g \psi_g \alpha_g}. \tag{17}$$

The expression for $\tau^R_{g,oF}$ in (15) shows that the government reduces the relative size of organizational modes in proportion to their conflict output loss, $\Delta_{g,oF}$. This loss is adjusted by the firm's ability to compensate for the loss of foreign headquarters inputs during conflict, $\frac{d \ln y^C_{g,oF}}{d \ln y_{g,oF}} = \frac{\eta^m_g}{1 - \eta^h_g \psi_g \alpha_g} \in (0, 1)$. When the adjusted loss is sufficiently large, it is optimal to set a positive revenue wedge for firms with foreign headquarters. More generally, revenue wedges are higher (i.e., imply smaller subsidies) for firms with greater reliance on foreign headquarters,

$$\tau^R_{g,VF} > \tau^R_{g,AF} > \tau^R_{g,VD} = \tau^R_{g,AD}.$$

Equation (16) characterizes the organizational wedge applied to firms selecting different organizational modes. This wedge forces firms to internalize how their organizational choice affects consumption in the off-equilibrium conflict state, as captured by $\lambda_g \Delta_{g,oF}$. Once these effects are accounted for, the output-loss component implies larger wedges (as a fraction of revenue) for firms that rely more intensively on foreign headquarters services,

$$\bar{\tau}^O_{g,VF} > \bar{\tau}^O_{g,AF} > \bar{\tau}^O_{g,VD} = \bar{\tau}^O_{g,AD}.$$

The intuition for the wedges in (15)–(16) is similar. The government seeks to reduce reliance on foreign headquarters services. Because it cannot directly contract on firms' input composition, it instead reduces the relative output of organizational modes whose non-contractible input composition relies more heavily on foreign headquarters services, at both the intensive and extensive margins through the revenue and organizational wedges, respectively.

Finally, the expression for $\tau^L_{g,oF}$ in (17) shows that, conditional on the optimal scale adjustment, it is optimal to increase the use of contractible labor in firms with foreign headquarters. This reduces their reliance on non-contractible inputs and therefore allows them to produce more in the conflict state.

Taken together, these wedges reallocate production away from organizational modes that are more exposed to geopolitical risk and, consequently, reduce the equilibrium international transfer in a crisis. They help to align the private and social values of relying on foreign headquarters services.

3.3 Determinants of Foreign Ownership Regulation

To build intuition for our main results, we present comparative statics that illustrate how different mechanisms shape government regulation of foreign firm ownership. We first study how regulation responds to an increase in geopolitical frictions. We then characterize how regulation varies across goods depending on their relevance for national defense and the importance of non-contractible inputs in production. Throughout, we rely on a calibration in which organizational modes are equally important in the absence of national security concerns.[11]

3.3.1 Geopolitical Risk

Our first set of comparative statics examines how geopolitical risk affects the optimal restrictiveness of foreign ownership. In the model, the parameter θ captures geopolitical risk through the probability of a crisis that triggers international bargaining between governments. It can be interpreted as the likelihood of geopolitical disputes that allow foreign countries to coerce the domestic country. As θ increases, we move from a relatively stable environment to one in which geopolitical crises are perceived to be more frequent.

Figure 1 shows how outcomes change as θ rises. All variables are normalized relative to the no-policy competitive equilibrium, which corresponds to the benchmark case of $\theta = 0$. As geopolitical risk increases, the expected international transfer rises. In response, the government strengthens its intervention by increasing the regulatory wedges it imposes on the private sector, as characterized in Appendix A.3. These wedges distort

[11] The baseline calibration sets $c^h_{g,\ell} = c^m_g = c^C_g = 1$ and $\eta^m_g = \eta^h_g = \frac{1}{2}$. It sets $\beta^h_V = 0.8$ and $\beta^h_A = 0.2$, so that $\beta^m_V = 0.2$ and $\beta^m_A = 0.8$. This symmetry implies that $P_{g,o\ell} = 0.25$, which is useful for exposition. The baseline further sets $\alpha_g = \psi_g = 0.75$, $\sigma_g = 5$, $v_g = 10$, $\theta = 0.3$, $d_g = 0.025$, and $d^C_g = 2d_g$.

allocations in the peace state to reduce the expected transfer in the event of a crisis. As a result, the economy becomes more resilient to conflict, which attenuates the increase in prices during conflict. Consistent with this mechanism, panel (a) of Figure 1 shows that λ_g declines with θ.

The last two panels of Figure 1 show how the government achieves the reduction in λ_g. The optimal wedges both increase the share of firms that source headquarters services domestically and expand their relative output. According to Proposition 2, the government does not intervene in the ownership structure of firms that source headquarters services domestically. Consequently, vertically integrated and arm's-length firms behave identically within this group. In contrast, Proposition 3 shows that the government targets firms that rely on foreign headquarters services by reducing both their frequency and output, reflecting their larger output losses during conflict. Restrictions are stronger for vertically integrated firms because their ownership structure induces greater reliance on foreign headquarters services.

Figure 1: Optimal Foreign Ownership Regulation and Geopolitical Risk

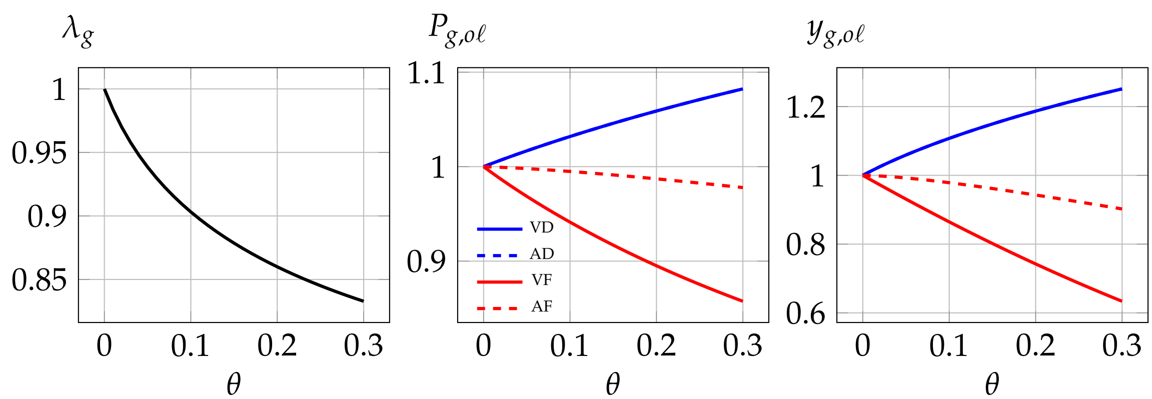

Note: The figure illustrates the impact of optimal policy on outcomes (relative to the competitive equilibrium) as a function of the probability of a geopolitical crisis, θ. Panel (a) reports the relative marginal value of consumption in conflict, λ_g. Panels (b) and (c) report the share and output of firms across organizational modes.

3.3.2 National Defense and Non-Contractible Inputs

We now turn to comparisons across sectors that differ in their importance for national defense, d_g, and their intensity of non-contractible inputs, α_g. We focus on a setting with high geopolitical risk, $\theta = 0.3$. Figure 2 displays how optimal wedges vary across ownership and location choices. It highlights three patterns that we investigate empirically in the next section.

First, as panel (a) shows, the greater the defense demand for a sector, the larger the restrictions the government places on foreign ownership. The government imposes larger wedges on foreign-owned firms relative to domestic ones as defense demand rises.[12] Appendix Figure 5 shows that this leads to a lower frequency of foreign-owned firms. Importantly, these wedges are larger for foreign vertically integrated firms than for foreign arm's-length firms, because the former rely more intensively on headquarters services and therefore experience larger output losses during conflict.

Second, as panel (b) shows, the greater the share of non-contractible inputs, the larger the wedges placed on foreign ownership. Appendix Figure 6 shows that this leads to a lower frequency of foreign-owned firms. When non-contractible inputs are more important, the output loss $\Delta_{g,oF}$ is larger because it becomes more costly to substitute toward the backup conflict technology when foreign headquarters services are disrupted. This mechanism amplifies the increase in the marginal value of consumption in conflict, λ_g. While the higher λ_g tends to reduce wedges overall, the effect is weaker for foreign firms, which suffer output losses, than for domestic firms, which do not. As a result, wedges are larger for foreign firms relative to domestic firms. Note that again the wedge gap is larger for vertically integrated firms.

Third, the stronger are both mechanisms, the tighter is regulation of foreign ownership. Regulation is particularly tight when goods are both important for national defense and rely heavily on non-contractible inputs. Panel (c) of Figure 2 illustrates this interaction through the steeper increase in the relative wedge for firms vertically integrated with foreign headquarters at higher levels of non-contractible input intensity. This arises because the impact of output losses on the marginal value of consumption increases with defense spending during conflict. When non-contractible inputs are more important, the cost of these output losses is therefore higher in sectors with greater defense demand. Consequently, regulation is strongest in sectors that combine high defense demand with high reliance on non-contractible inputs.

3.4 Extensions

We now turn to extensions. We give a brief summary in the main text and provide details in Appendix A.2

[12] By Proposition 2, we have $\bar{\tau}^O_{g,VD} = \bar{\tau}^O_{g,AD}$, so the choice of domestic organizational form does not affect the benchmark.

Figure 2: Optimal Foreign Ownership Regulation Across Goods

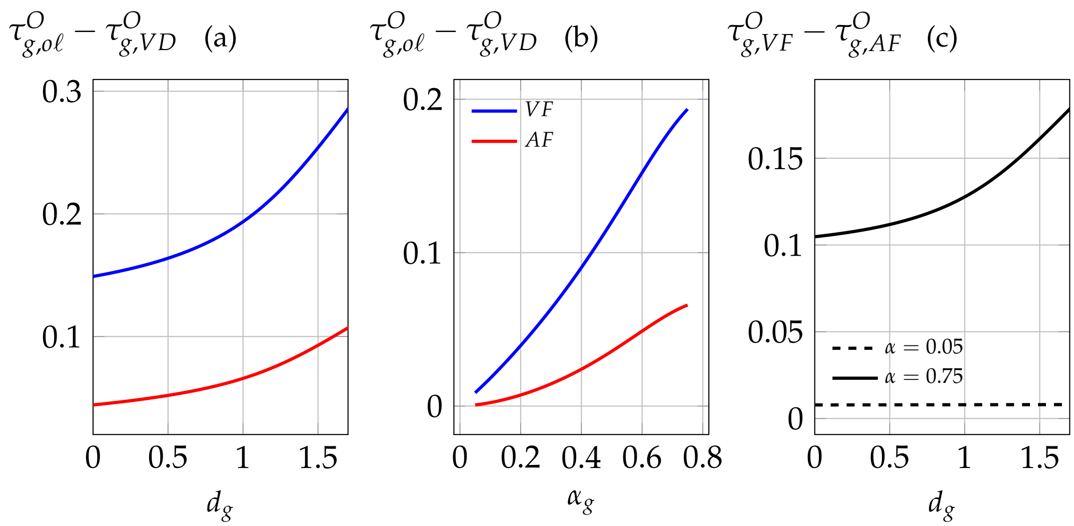

Note: The figure illustrates the impact of optimal policy on outcomes across goods that differ in their defense demand d_g and the share of non-contractible inputs (α_g). Reported defense demands are normalized by the baseline defense demand. Panels (a) and (b) report the wedges of different organizational modes, relative to firms vertically integrated with domestic headquarters. Panel (c) reports the wedge of foreign vertically integrated firms (VF) relative to foreign arm's-length firms (AF). Appendix A.3 reports additional figures with wedges and firm-level changes in output and input composition.

Implementation through taxes. The baseline model decentralizes the planner's allocation through the use of regulatory wedges. In the extension we discuss how outcomes are affected if each of our three wedges is replaced by a corresponding tax. The main result is that revenue and labor taxes have identical expressions to the ones in the main text. The expressions for the direct national security benefits of location and ownership choices are also unchanged. But the ownership and location tax picks up some additional terms compared to its wedge counterpart. The reason is that when the government uses revenue and labor taxes differentially across ownership structures and locations, this distorts firms' decisions towards those choices that are encouraged. The government uses the ownership and location tax to offset these effects and replicate the revenue-neutral wedges in our baseline model.

4 Foreign Ownership Regulation in Practice: The Case of the United States

In this section, we confront the model's predictions with data on how the United States regulates foreign ownership and control of domestic firms. We first describe the institutional setting through which the U.S. government regulates foreign acquisitions, and how the U.S. regulatory regime changed with the evolution of geopolitical risks over time. We then relate cross-sector variation in regulatory intensity to observable proxies for the theoretical determinants of optimal intervention.

4.1 Institutional Background

The Committee on Foreign Investment in the United States (CFIUS) serves as the primary institution through which the U.S. government regulates foreign ownership and control of domestic firms. It is an interagency committee, chaired by the Secretary of the Treasury, that reviews foreign investment in U.S. businesses in order to identify and mitigate risks to national security. When national security risks are identified, CFIUS may negotiate mitigation agreements with the transacting parties or, in the most serious cases, recommend that the President block or unwind the transaction.

Legislative evolution. The evolution of CFIUS over time has been shaped by the impact of changes in the geopolitical landscape on perceptions of national security risks posed by foreign ownership and control of U.S. firms.

CFIUS was established by Executive Order in 1975 as an advisory body in response to concerns over investments by OPEC members in U.S. firms.[13] In 1988, Japanese acquisitions of defense-relevant firms led Congress to give the President standing authority to block foreign acquisitions that threaten national security.[14] CFIUS jurisdiction was to review transactions that confer control of a U.S. business to a foreign person. In an era of relative geopolitical stability, this scope remained largely unchanged even as foreign direct investment in the United States grew from $395 billion in 1990 to $3.4 trillion in 2015.

The emergence of China as a geopolitical competitor triggered updates to the regulatory framework governing foreign investment in U.S. firms. As geopolitical risks increased, policymakers recognized that foreign investors could take actions within U.S. firms that the existing regulatory framework was not designed to monitor.[15] The Foreign

[13] CRS Report RL33388, *The Committee on Foreign Investment in the United States*, by James K. Jackson.

[14] The Exon-Florio Amendment of 1988 codified this expansion of powers. The most prominent case was Fujitsu's attempted purchase of Fairchild Semiconductor, a military chip supplier ("Cold Feet: Fujitsu Drops Its Fairchild Bid," *TIME*, March 30, 1987; CRS Report RL33388). The Byrd Amendment of 1992 mandated investigations of acquisitions by foreign government-controlled entities (CRS Report IF10177).

[15] In 2007, the Foreign Investment and National Security Act (FINSA) codified CFIUS in statute and strengthened Congressional oversight. FINSA was prompted by the withdrawal of CNOOC's bid for Unocal under Congressional pressure (2005) and the Dubai Ports World controversy (2006). See "Chinese Drop Bid To Buy U.S. Oil Firm," *The Washington Post*, August 3, 2005; CRS Report RL33388.

Investment Risk Review Modernization Act (FIRRMA) of 2018 marked a sharp expansion of CFIUS jurisdiction. A central concern was that, without triggering CFIUS review, non-controlling investment structures—such as minority stakes, joint ventures, and licensing arrangements—allowed foreign investors to shape decisions within firms, access managerial know-how, and facilitate technology transfers. Congressional hearings and FIRRMA's statutory text explicitly linked these concerns to rising strategic competition with China and its Made in China 2025 industrial policy.[16]

After its implementation in February 2020, FIRRMA expanded CFIUS jurisdiction beyond control transactions to include non-controlling foreign investments across all sectors of the economy. It introduced mandatory filing requirements, an expedited declaration pathway, and review authority over real estate transactions near sensitive government facilities. FIRRMA also established heightened scrutiny over transactions involving critical technologies and infrastructure, as well as sensitive personal data (TID U.S. businesses).

The trajectory of these reforms is consistent with the model's prediction that rising geopolitical risk θ leads to more intensive regulation of foreign ownership. As strategic competition intensified, firms' organizational decisions increasingly determined their reliance on foreign inputs in ways that regulators could not directly observe or control. Ownership regulation therefore became the government's primary instrument to influence that reliance.

Review process. Under the current regime, parties can submit a filing to CFIUS through one of two parallel pathways. The first is a *declaration*: a short-form filing that triggers a 30-day assessment. CFIUS may clear the transaction, request that the parties file a full notice, or conclude that it is unable to complete its assessment. The second is a *notice*: a comprehensive submission that initiates a 45-day review, potentially followed by a 45-day investigation and a 15-day Presidential decision period. Whenever CFIUS concludes that a notice poses national security risks, it may propose mitigation agreements that impose conditions on the transaction.

In recent years, CFIUS has reviewed a substantial and growing volume of transactions: in 2024, 209 notices and 116 declarations were filed.[17] In contrast, CFIUS reviewed on average only 125 notices per year between 2010 and 2016, before the enactment of FIR-

[16]For example, see the testimony of Senator Cornyn in the Senate Committee on Banking, Housing, and Urban Affairs, hearing on CFIUS reform, January 18, 2018. FIRRMA's statutory text encodes this directly: Section 1702(c) requires the Secretary of Commerce to report to Congress on Chinese investment patterns and their alignment with the goals of Made in China 2025 (P.L. 115-232, § 1702(c)).

[17]In 2024, 16 of 209 notices were concluded with mitigation agreements, 49 were withdrawn (of which 42 were subsequently refiled), and 7 were abandoned. Presidential decisions to block transactions remain rare: only approximately ten have been issued since 1990.

RMA. Appendix Figure 7 shows that the number of notices and declarations increased substantially between 2005 and 2023. Appendix Figure 8 further shows that the share of notices concluded with mitigation agreements increased from 2% in 2008 to 15% in 2023.

Filing is voluntary for most transactions: parties choose to notify CFIUS in order to obtain safe harbor protection, which prevents the committee from later reopening a review about a concluded transaction. Filing is mandatory only when the transaction involves critical technologies requiring export authorization to the foreign acquirer, or when a foreign government holds a substantial interest in the acquiring entity and the target is a TID U.S. business. In 2024, only 31% of declarations followed mandatory requirements.[18]

4.2 Data

We now describe the construction of the variables used in our empirical analysis, based on variation in regulatory intensity across 4-digit NAICS manufacturing industries, which we map to goods g in the model.

Regulatory intensity. We use the CFIUS Annual Reports to Congress to obtain the number of notices and declarations filed each year by the 4-digit NAICS industry of the U.S. business under review. These tabulations provide the finest publicly available information on the sectoral distribution of CFIUS reviews, as transaction-level data are statutorily exempt from disclosure for privacy reasons. We harmonize all industry codes to the 2017 NAICS classification.

To proxy for the regulatory wedges in our model, we compute the total number of filings in each sector, summing notices and declarations, which capture variation in the intensity of foreign ownership regulation across industries. This aggregation reflects the fact that parties may satisfy filing requirements through either pathway, and that the choice between them depends on transaction-level characteristics rather than sector-level attributes. We aggregate filings over the period 2021–2024, which ensures a consistent institutional environment throughout the sample, as FIRRMA was fully implemented in 2020. We verify that our results are robust to using each filing type separately and to including filings from the pre-FIRRMA period.

The resulting measure exhibits substantial cross-sector variation. Appendix Figure 9 reports number of filings by two-digit NAICS sector. Manufacturing (NAICS 31-33) alone accounts for nearly a third of all filings in our estimation sample, followed by Professional & Technical Services (NAICS 54) and Information (NAICS 51). Appendix Table

[18] Annual reports do not distinguish mandatory from voluntary notices.

2 disaggregates the top ten four-digit industries within manufacturing. Filings are concentrated in semiconductors, aerospace, and communications equipment. This variation forms the basis of our empirical strategy: we relate the cross-sector distribution of filings to observable proxies for the theoretical determinants of optimal intervention.

Defense demand, d_g. To proxy for defense demand in the model, we measure each sector's exposure to Department of Defense procurement. We obtain contract-level data from the Federal Procurement Data System (FPDS) via USAspending.gov for fiscal years 2021–2024 and classify contracts as defense-related based on the awarding agency. To account for indirect supply-chain linkages, we scale direct procurement using the Leontief inverse derived from the 2017 BEA Input-Output Use Table. The resulting measure captures each industry's total contribution to defense demand, including both direct procurement and indirect demand through intermediate input linkages.

Non-contractible input intensity, α_g. To proxy for non-contractible input intensity in the model, we use the contract intensity index of Nunn (2007). This index measures, for each industry, the share of intermediate inputs that are relationship-specific, based on the classification of goods in Rauch (1999) as traded on organized exchanges, reference-priced, or differentiated. We use the original index constructed from the 1997 U.S. Input-Output Table and map it to four-digit NAICS 2017 industries, weighting by input values when multiple codes correspond to the same NAICS industry. The dataset covers 115 industries, of which 85 are in manufacturing. The availability of this measure motivates our focus on manufacturing industries, although the results are similar when we extend the analysis to non-manufacturing industries with available data.[19]

4.3 Foreign Ownership Regulatory Intensity Across Manufacturing Industries in the United States

We begin by replicating the comparative statics exercise in Section 3.3 using our dataset on regulatory intensity in the United States. Figure 3 reports binscatter plots illustrating the relationship across manufacturing industries between regulatory intensity and two key determinants of optimal foreign ownership restrictions in the model: defense demand

[19] We obtain two additional measures of non-contractibility intensity for manufacturing industries used in the literature. Specifically, following Antràs (2003), we use the capital-labor ratio from the NBER Manufacturing Database, and following Levchenko (2007), we use the negative of the Herfindahl-Hirschman Index (HHI) of input spending shares.

(d_g) and non-contractible input intensity (α_g). These figures provide a non-parametric empirical counterpart to the model's core predictions.

Panel (a) of Figure 3 shows that regulatory intensity increases with defense demand. The convex relationship aligns closely with the model's prediction in panel (a) of Figure 2. This suggests that restrictions on foreign ownership are particularly acute in industries that are central to national defense. In the model, this reflects their importance for the country's bargaining power during geopolitical crises.

Panel (b) of Figure 3 examines the relationship between regulatory intensity and non-contractible inputs across industries. Consistent with the model's predictions in panel (b) of Figure 2, the number of filings increases with the industry's reliance on non-contractible inputs.

Finally, panel (c) of Figure 3 shows that the relationship between regulatory intensity and defense demand is driven primarily by industries that rely heavily on non-contractible inputs. This pattern is consistent with panel (c) of Figure 2 and illustrates that foreign ownership regulation is particularly strong in industries that are both important for national defense and intensive in non-contractible inputs.

Figure 3: Regulatory Intensity, Defense Demand, and Non-Contractibility Input Intensity

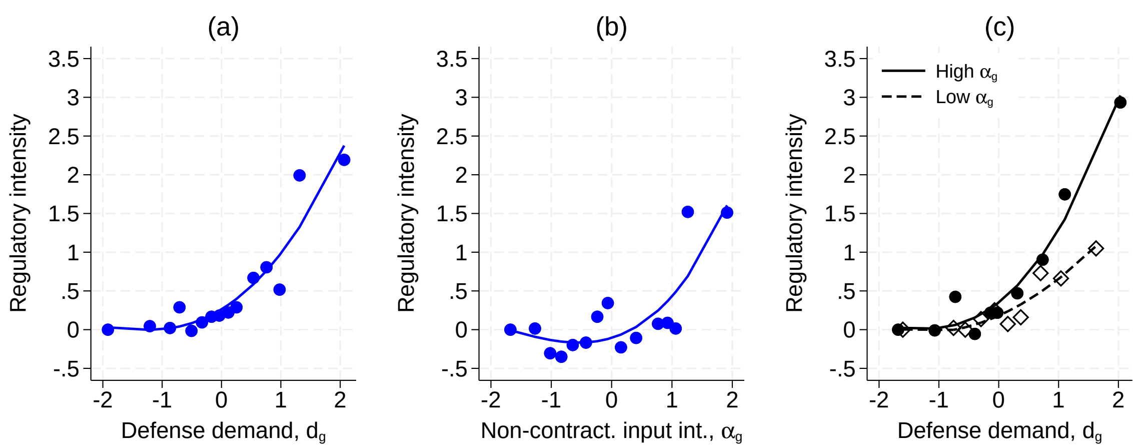

Notes: Each panel plots binned means of regulatory intensity against industry characteristics for a sample of 85 four-digit NAICS manufacturing industries. Fitted lines are fractional polynomial fits to the binned means. Panel (a) plots regulatory intensity against defense demand (d_g). Panel (b) plots regulatory intensity against non-contractible input intensity (α_g). Panel (c) splits the sample into above- and below-mean α_g and plots regulatory intensity against defense demand separately for each group. All variables are defined in Section 4.2 and are standardized to have unit standard deviation. Each series is normalized so that regulatory intensity equals zero in the lowest bin of the corresponding industry characteristic.

We now turn to a formal regression analysis of these relationships across industries.

In particular, we estimate the following linear regression:

$$\text{Regulatory intensity}_g = \beta^d d_g + \beta^\alpha \alpha_g + \beta^{\alpha d} d_g \alpha_g + \text{Controls}_g + \epsilon_g \tag{18}$$

where all variables are defined in Section 4.2, and Controls_g is a vector of control variables. The model predicts positive values for β^d, β^α, and $\beta^{\alpha d}$. These coefficients directly correspond to the comparative statics in Section 3.3.

Table 1 reports the estimates of specification (18). Column (1) shows that the univariate relationships in Figure 3 also hold jointly in our sample. Defense demand and non-contractible input intensity together explain 53% of the variation in regulatory intensity across the 85 four-digit manufacturing industries in our sample.

The estimated coefficients are economically meaningful. Because all variables are standardized, the coefficients can be interpreted as the change in regulatory intensity (in standard deviations) associated with a one standard deviation increase in each explanatory variable. In particular, a one standard deviation increase in defense demand is associated with an increase of 0.469 standard deviations in regulatory intensity for a sector with the mean level of non-contractible input intensity. Similarly, a one standard deviation increase in non-contractible input intensity raises regulatory intensity by 0.103 standard deviations for a sector with the mean level of defense demand. The positive interaction term implies that these effects are amplified in sectors that are both strategically important and intensive in non-contractible inputs, consistent with the model's predictions.

The remaining columns of Table 1 evaluate the robustness of these findings. Columns (2) and (3) show that the estimates are similar when we consider the two filing paths separately, declarations and notices. Column (4) reports estimates for notices in the pre-FIRRMA period (2014–2020), when notices were the only filing path. The estimates are similar, although less precise.

Column (5) considers additional determinants of filings. We include two measures of industry size—the log of output and the log of number of firms—which may mechanically increase the likelihood of review. We also include two-digit NAICS fixed effects to capture broad industry characteristics. These controls have little effect on the estimates.

In column (6), we show that the results are similar, though less precise, when we expand the sample to include non-manufacturing industries with available data on non-contractible input intensity. This extended sample includes sectors that are relevant for foreign ownership regulation, such as utilities and software publishing.

Finally, column (7) shows that our main findings are robust to controlling for alternative measures of non-contractible input intensity used in the literature and their interactions with defense demand. In particular, we consider the capital-labor ratio, as in Antràs (2003), and the dispersion in input spending, as in Levchenko (2007). The baseline coefficients remain largely unchanged when including these controls, despite their additional explanatory power. Appendix Table 3 reports the coefficients on these variables, which also have signs consistent with our theoretical predictions.

Table 1: CFIUS Regulatory Intensity and Sector Characteristics

Dep. var.:	All filings Baseline	Decl. 2021–24	Notices 2021–24	Notices 2014–20	All filings Robustness			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Defense demand (d_g)	0.505	0.487	0.503	0.409	0.449	0.453	0.437	0.118
	(0.090)	(0.096)	(0.089)	(0.096)	(0.100)	(0.124)	(0.091)	(0.048)
Non-contract. input int. (α_g)	0.275	0.284	0.262	0.254	0.282	0.346	0.284	0.253
	(0.076)	(0.071)	(0.081)	(0.090)	(0.078)	(0.091)	(0.078)	(0.081)
$d_g \times \alpha_g$	0.265	0.289	0.243	0.200	0.267	0.349	0.265	0.118
	(0.079)	(0.089)	(0.075)	(0.075)	(0.079)	(0.110)	(0.072)	(0.060)
R^2	0.529	0.541	0.495	0.353	0.542	0.638	0.545	0.517
Observations	85	85	85	85	85	85	108	292
Sample	Manuf.	Manuf.	Manuf.	Manuf.	Manuf.	Manuf.	Tradables	All sectors
Size controls	No	No	No	No	Yes	Yes	Yes	Yes
Alt. contract int. measures	No	No	No	No	No	Yes	No	No
Sector group dummies	No	No	No	No	No	No	Yes	Yes

Notes: Each column reports OLS estimates of (18) across four-digit NAICS industries. The dependent variable is the standardized number of CFIUS filings. Column (2) uses declarations only (2021–2024), column (3) uses notices only (2021–2024, post-FIRRMA), and column (4) uses notices only (2014–2020, pre-FIRRMA). Column (5) includes controls for log total output and for the number of inbound cross-border M&A deals targeting U.S. firms during 2021–2024. Column (6) additionally controls for the capital-labor ratio following Antràs (2003) and the negative of the Herfindahl-Hirschman Index of input spending shares following Levchenko (2007), each interacted with defense demand; coefficients are reported in Appendix Table 3. The sample in columns (1)-(6) includes 85 manufacturing industries. Column (7) expands the sample to all tradable sectors in agriculture (11), mining (21) and manufacturing (31-33), and column (8) to all four-digit NAICS industries; both control for log total output, the number of inbound cross-border M&A deals, and a set of sector group dummies—primary (11-21), utilities (22), manufacturing (31-33), transportation & trade (42-45), information (51), business services (52-56). All variables are defined in Section 4.2 and standardized using the manufacturing sample moments. Robust standard errors are reported in parentheses.

5 Conclusion

This paper develops a theory of foreign ownership regulation based on the interaction of incomplete contracts within firms and geopolitical bargaining between governments. In our framework, ownership structure and headquarters location shape firms' reliance on foreign non-contractible inputs. While these choices are privately optimal given the contractual environment, firms do not internalize how they affect the domestic government's outside option in a geopolitical crisis. This creates a national security externality. Because the government is subject to the same contractual limitations as private agents, it cannot directly regulate the underlying non-contractible margin. Foreign ownership regulation therefore emerges as the optimal second-best response: by shifting firms away from organizational forms that rely more heavily on foreign non-contractible inputs, it improves resilience in crisis and strengthens national bargaining power.

The theory also delivers sharp predictions about the cross-sector pattern of intervention. Regulation should be stricter in sectors that are more important in conflict and in sectors that rely more heavily on non-contractible inputs, with particularly strong intervention where both features coincide. Using a new dataset on CFIUS reviews, we show that these predictions are borne out in the data: sectors with greater military exposure and higher non-contractibility face tighter foreign ownership regulation, and together these forces explain a substantial share of observed cross-industry variation. More generally, the paper highlights that when the relevant source of geopolitical vulnerability is embedded in firms' organizational choices, ownership regulation can play a central role in optimal national security policy.

References

Alekseev, Maxim and Xinyue Lin, "Export Policy for Dual-Use Goods," Technical Report, Working paper February 2026. First version circulated on November 1, 2025, under the title "Trade Policy in the Shadow of Conflict: The Case of Dual-Use Goods".

Antràs, Pol, "Firms, contracts, and trade structure," *The Quarterly Journal of Economics*, 2003, *118* (4), 1375–1418.

_ **and Elhanan Helpman**, "Global sourcing," *Journal of political Economy*, 2004, *112* (3), 552–580.

Becko, John Sturm and Daniel O'Connor, "Strategic (dis) integration," Technical Report, Working paper 2024.

Boehm, Christoph E, Aaron Flaaen, and Nitya Pandalai-Nayar, "Input linkages and the transmission of shocks: Firm-level evidence from the 2011 Tōhoku earthquake," *Review of Economics and Statistics*, 2019, *101* (1), 60–75.

Clayton, Christopher, Matteo Maggiori, and Jesse Schreger, "A framework for geoeconomics," Technical Report, National Bureau of Economic Research 2023.

_ **,** _ **, and** _ , "A theory of economic coercion and fragmentation," *Available at SSRN 4767131*, 2024.

Costinot, Arnaud, Guido Lorenzoni, and Iván Werning, "A theory of capital controls as dynamic terms-of-trade manipulation," *Journal of Political Economy*, 2014, *122* (1), 77–128.

Farhi, Emmanuel and Iván Werning, "A theory of macroprudential policies in the presence of nominal rigidities," *Econometrica*, 2016, *84* (5), 1645–1704.

Gilpin, Robert, *U.S. Power and the Multinational Corporation: The Political Economy of Foreign Direct Investment*, New York: Basic Books, 1975.

Grossman, Sanford J and Oliver D Hart, "The costs and benefits of ownership: A theory of vertical and lateral integration," *Journal of political economy*, 1986, *94* (4), 691–719.

Hart, Oliver and John Moore, "Property Rights and the Nature of the Firm," *Journal of political economy*, 1990, *98* (6), 1119–1158.

_ **and** _ , "Foundations of incomplete contracts," *The Review of Economic Studies*, 1999, *66* (1), 115–138.

Kooi, Oliver, "Power and resilience: An economic approach to national security policy," Technical Report, working paper 2024.

Levchenko, Andrei A, "Institutional quality and international trade," *The Review of Economic Studies*, 2007, *74* (3), 791–819.

McLaren, John, "Size, sunk costs, and Judge Bowker's objection to free trade," *The American Economic Review*, 1997, pp. 400–420.

Mohr, Cathrin and Christoph Trebesch, "Geoeconomics," *Annual Review of Economics*, 2025, *17*.

Nunn, Nathan, "Relationship-specificity, incomplete contracts, and the pattern of trade," *The quarterly journal of economics*, 2007, *122* (2), 569–600.

Rauch, James E, "Networks versus markets in international trade," *Journal of international Economics*, 1999, *48* (1), 7–35.

Segal, Ilya, "Complexity and renegotiation: A foundation for incomplete contracts," *The Review of Economic Studies*, 1999, *66* (1), 57–82.

Waltz, Kenneth N., *Man, the State, and War: A Theoretical Analysis*, New York, NY: Columbia University Press, 1959.

Wellhausen, Rachel L, *The shield of nationality*, Cambridge University Press, 2015.

A Theoretical Appendix

A.1 Proofs of Propositions 1–3

We now establish the proof for Propositions 1–3. Note that Propositions 1 is the special case with $\theta = 0$.

The full problem of the government is given by

$$\max_{\{P_{g,o\ell}, y_{g,o\ell}, l_{g,o\ell}, h_{g,o\ell}, m_{g,o\ell}\}} \mathcal{U} + \theta \sum_{g=1}^{G} \frac{\sigma_g}{\sigma_g - 1} \left((y_g^C - d_g^C)^{\frac{\sigma_g - 1}{\sigma_g}} - (y_g - d_g)^{\frac{\sigma_g - 1}{\sigma_g}} \right) \quad \text{subject to}$$

$$\sum_{o,\ell} P_{g,o\ell} = 1, \quad h_{g,o\ell} = \frac{\beta_{g,o}^h \eta_g^h c_g^m}{\beta_{g,o}^m \eta_g^m c_{g,\ell}^h} m_{g,o\ell}, \quad y_{g,o\ell} = \left[(l_{g,o\ell})^{1-\alpha_g} \left((h_{g,o\ell})^{\eta_g^h} (m_{g,o\ell})^{\eta_g^m} \right)^{\alpha_g} \right]^{\psi_g},$$

$$y_{g,oD}^C = y_{g,oD}, \quad y_{g,oF}^C = \rho_g^C \left((m_{g,oF})^{\eta_g^m \alpha_g} (l_{g,oF})^{1-\alpha_g} \right)^{\frac{\psi_g}{1-\eta_g^h \alpha_g \psi_g}}.$$

Simplifying. It is useful to allow the planner to directly choose contractible inputs $l_{g,o\ell}$ and the non-contractible bundle $X_{g,o\ell} \equiv h_{g,o\ell}^{\eta_g^h} m_{g,o\ell}^{\eta_g^m}$. Define the equilibrium input ratio

$$\phi_{g,o\ell} \equiv \frac{\beta_{g,o}^h \, \eta_g^h \, c_g^m}{\beta_{g,o}^m \, \eta_g^m \, c_{g,\ell}^h},$$

which is determined solely by parameters and follows directly from the non-contractibility constraint $h_{g,o\ell} = \phi_{g,o\ell} \, m_{g,o\ell}$. Given $\phi_{g,o\ell}$ and $X_{g,o\ell}$, the individual inputs are uniquely determined:

$$h_{g,o\ell} = \phi_{g,o\ell}^{\eta_g^m} X_{g,o\ell}, \qquad m_{g,o\ell} = \phi_{g,o\ell}^{-\eta_g^h} X_{g,o\ell}.$$

The associated unit cost of the bundle $X_{g,o\ell}$ is $c_{g,o\ell}^X = c_{g,\ell}^h \, \phi_{g,o\ell}^{\eta_g^m} + c_g^m \, \phi_{g,o\ell}^{-\eta_g^h}$.

The planning problem can therefore be restated as

$$
\begin{aligned}
\max_{\{P_{g,o\ell}, y_{g,o\ell}, l_{g,o\ell}, X_{g,o\ell}\}} \quad & L + \sum_{g=1}^{G} \frac{\sigma_g}{\sigma_g - 1} (y_g - d_g)^{\frac{\sigma_g - 1}{\sigma_g}} - \sum_{g,o,\ell} P_{g,o\ell} \left(c^X_{g,o\ell}\, X_{g,o\ell} + l_{g,o\ell} + \nu_g^{-1} \log P_{g,o\ell} \right) \\
& + \theta \sum_{g=1}^{G} \frac{\sigma_g}{\sigma_g - 1} \left((y^C_g - d^C_g)^{\frac{\sigma_g - 1}{\sigma_g}} - (y_g - d_g)^{\frac{\sigma_g - 1}{\sigma_g}} \right) \\
\text{subject to} \quad & \sum_{o,\ell} P_{g,o\ell} = 1, \qquad y_{g,o\ell} = \left[(l_{g,o\ell})^{1-\alpha_g} \left(X_{g,o\ell} \right)^{\alpha_g} \right]^{\psi_g}, \\
& y^C_{g,oD} = y_{g,oD}, \qquad y^C_{g,oF} = \rho^C_g \left(\left(\phi_{g,oF}^{-\eta^h_g} X_{g,oF} \right)^{\eta^m_g \alpha_g} \left(l_{g,oF} \right)^{1-\alpha_g} \right)^{\frac{\psi_g}{1 - \eta^h_g \alpha_g \psi_g}}.
\end{aligned}
$$

Optimal subsidies for contractible inputs The first order condition for the planner with respect to $l_{g,o\ell}$ is

$$
(q_g)^{-\frac{1}{\sigma_g}} \frac{\partial y_{g,o\ell}}{\partial l_{g,o\ell}} + \theta \left[\left(q^C_g \right)^{-\frac{1}{\sigma_g}} \frac{\partial y^C_{g,o\ell}}{\partial l_{g,o\ell}} - (C_g)^{-\frac{1}{\sigma_g}} \frac{\partial y_{g,o\ell}}{\partial l_{g,o\ell}} \right] - 1 = 0 \tag{19}
$$

The corresponding first order condition for a firm is

$$
p_g \frac{\partial y_{g,o\ell}}{\partial l_{g,o\ell}} (1 - \tau^R_{g,o\ell}) - (1 + \tau^L_{g,o\ell}) = 0 \tag{20}
$$

Combining the two equations and using that $\left(C^Z_g \right)^{-\frac{1}{\sigma_g}} = p^Z_g$ one obtains

$$
\tau^L_{g,o\ell} = -p_g \frac{\partial y_{g,o\ell}}{\partial l_{g,o\ell}} \left(\tau^R_{g,o\ell} + \theta \left[\frac{p^C_g y^C_{g,o\ell} \frac{\partial \ln y^C_{g,o\ell}}{\partial l_{g,o\ell}}}{p_g y_{g,o\ell} \frac{\partial \ln y_{g,o\ell}}{\partial l_{g,o\ell}}} - 1 \right] \right) \tag{21}
$$

Under the assumed technologies we know that $\frac{\partial y_{g,o\ell}}{\partial l_{g,o\ell}} = \psi_g (1 - \alpha_g) \frac{y_{g,o\ell}}{l_{g,o\ell}}$, $\frac{\partial y^C_{g,oF}}{\partial l_{g,oF}} = \frac{(1-\alpha_g)\psi_g}{1 - \eta^h_g \alpha_g \psi_g} \frac{y^C_{g,oF}}{l_{g,oF}}$, and $\frac{\partial y^C_{g,oD}}{\partial l_{g,oD}} = (1 - \alpha_g) \psi_g \frac{y^C_{g,oD}}{l_{g,oD}}$. Hence the expression simplifies to

$$
\tau^L_{g,o\ell} = -\frac{p_{g,o\ell} y_{g,o\ell}}{l_{g,o\ell}} \psi_g (1 - \alpha_g) \left(\tau^R_{g,o\ell} + \theta \left[\frac{p^C_g y^C_{g,o\ell}}{p_g y_{g,o\ell}} \Psi^l_{g,\ell} - 1 \right] \right) \tag{22}
$$

where $\Psi^l_{g,D} \equiv 1$ and $\Psi^l_{g,F} \equiv \frac{1}{1 - \eta^h_g \alpha_g \psi_g}$.

Optimal scale subsidies The first order condition of the planner with respect to $X_{g,o\ell}$ is

$$\left(C_g\right)^{-\frac{1}{\sigma_g}} \frac{\partial y_{g,o\ell}}{\partial X_{g,o\ell}} - c^X_{g,o\ell} + \theta \left[\left(C^C_g\right)^{-\frac{1}{\sigma_g}} \frac{\partial y^C_{g,o\ell}}{\partial X_{g,o\ell}} - \left(C_g\right)^{-\frac{1}{\sigma_g}} \frac{\partial y_{g,o\ell}}{\partial X_{g,o\ell}} \right] = 0 \tag{23}$$

The optimality condition for the firm is

$$p_g \frac{\partial y_{g,o\ell}}{\partial X_{g,o\ell}} (1 - \tau^R_{g,o\ell}) - c^X_{g,o\ell} = 0 \tag{24}$$

Combining the two equations and using that $\left(C^Z_g\right)^{-\frac{1}{\sigma_g}} = p^Z_g$ one obtains

$$\tau^R_{g,o\ell} = \theta \left[1 - \frac{p^C_g \frac{\partial y^C_{g,o\ell}}{\partial X_{g,o\ell}}}{p_g \frac{\partial y_{g,o\ell}}{\partial X_{g,o\ell}}} \right] \tag{25}$$

Under the assumed technologies we know that $\frac{\partial y_{g,o\ell}}{\partial X_{g,o\ell}} = \psi_g \alpha_g \frac{y_{g,o\ell}}{X_{g,o\ell}} = \psi_g \alpha_g \frac{y_{g,o\ell}}{m_{g,o\ell}} \phi^{-\eta^h_g}_{g,o\ell}$ and $\frac{\partial y^C_{g,oF}}{\partial X_{g,oF}} = \frac{\partial y^C_{g,oF}}{\partial m_{g,oF}} \frac{\partial m_{g,oF}}{\partial X_{g,oF}} = \frac{\alpha_g \psi_g \eta^m_g}{1 - \eta^h_g \alpha_g \psi_g} \frac{y^C_{g,oF}}{m_{g,oF}} \phi^{-\eta^h_g}_{g,o\ell}$. This implies that the expression simplifies to

$$\tau^R_{g,o\ell} = \theta \left[1 - \frac{p^C_g y^C_{g,o\ell}}{p_g y_{g,o\ell}} \Psi_{g,\ell} \right] \tag{26}$$

where $\Psi_{g,D} \equiv 1$ and $\Psi_{g,F} \equiv \frac{\eta^m_g}{1 - \eta^h_g \alpha_g \psi_g}$.

Optimal Location Decision The first order condition with respect to $P_{g,o\ell}$ is

$$\left(C_g\right)^{-\frac{1}{\sigma_g}} y_{g,o\ell} - c^h_{g,\ell} h_{g,o\ell} - c^m_g m_{g,o\ell} - l_{g,o\ell} + \theta \left[\left(C^C_g\right)^{-\frac{1}{\sigma_g}} y^C_{g,o\ell} - \left(C_g\right)^{-\frac{1}{\sigma_g}} y_{g,o\ell} \right] - \nu^{-1}_g \left(\ln P_{g,o\ell} + 1\right) = \mu_g \tag{27}$$

where μ_g is the multiplier on $\sum_{o,\ell} P_{g,o\ell} = 1$.

The Gumbel distribution implies the logit choice formula.[20] Hence

$$\nu^{-1}_g \left(\ln P_{g,o\ell} + 1\right) = \pi_{g,o\ell} + \bar{F}_{g,o\ell} - \tau^O_{g,o\ell} - N_g \tag{28}$$

[20] That is, $P_{g,o\ell} = \frac{\exp\left(\nu_g (\pi_{g,o\ell} + \bar{F}_{g,o\ell} - \tau^O_{g,o\ell})\right)}{\sum_{o',\ell'} \exp\left(\nu_g (\pi_{g,o'\ell'} + \bar{F}_{g,o'\ell'} - \tau^O_{g,o'\ell'})\right)}$.

where

$$N_g \equiv \nu_g^{-1} \ln \sum_{o',\ell'} \exp\Big(\nu_g(\pi_{g,o'\ell'} + \bar{F}_{g,o'\ell'} - \tau^O_{g,o'\ell'})\Big) - \nu_g^{-1}. \tag{29}$$

Using that $\left(C_g^Z\right)^{-\frac{1}{\sigma_g}} = p_g^Z$ and the expression for profits we obtain

$$\tau^O_{g,o\ell} = -\theta \left[p_g^C y^C_{g,o\ell} - p_g y_{g,o\ell}\right] + \mu_g - N_g \tag{30}$$

Scaling the $\tau^O_{g,o\ell}$ We can use the well-known location invariance of the Gumbel/logit formula: choice probabilities $P_{g,o\ell}$ are unchanged if we add the same constant to all wedges within good g. In particular, for any constant A_g, the wedge vectors $\tau^O_{g,o\ell}$ and $\tilde{\tau}^O_{g,o\ell} \equiv \tau^O_{g,o\ell} + A_g$ generate exactly the same probabilities $P_{g,o\ell}$, because only relative payoffs matter in the logit expression. We can therefore choose the normalization of $\tau^O_{g,o\ell}$ conveniently. In particular, we may pick A_g so that $N_g = \mu_g$.

A.2 Proofs for the Extensions

A.2.1 Implementation with taxes

Setup This appendix examines how the allocations in the baseline model can be implemented using taxes rather than regulatory wedges. Let $\{\tau^R_{g,o\ell}, \tau^L_{g,o\ell}, \tau^O_{g,o\ell}\}$ denote taxes that raise revenue for the government. Relative to the baseline model, this changes two things. First, lump-sum taxes are now used to rebate the revenue collected through these instruments. Second, firm profits are given by

$$\pi_{g,o\ell} = (1 - \tau^R_{g,o\ell}) p_g y_{g,o\ell} - c^h_{g,\ell} h_{g,o\ell} - c^m_g m_{g,o\ell} - (1 + \tau^L_{g,o\ell}) l_{g,o\ell}. \tag{31}$$

As a result, firms internalize the revenues or costs associated with location-ownership taxes when making those choices.

Optimal Policy The planner's optimality conditions for contractible inputs and for the bundle of non-contractible inputs are unchanged. The private-sector optimality conditions are likewise unchanged. The optimal revenue and labor taxes are therefore still given by Propositions 2 and 3.

The planner's optimality condition for location-ownership choice is also unchanged, but firms now take into account the revenues associated with the government's location-ownership taxes. Combining 27 with the profit expression above and the expression for the shares yields

$$\tau^O_{g,o\ell} = -p_g y_{g,o\ell} \tau^R_{g,o\ell} - l_{g,o\ell} \tau^L_{g,o\ell} - \theta \left[p^C_g y^C_{g,o\ell} - p_g y_{g,o\ell} \right]. \quad (32)$$

The first two terms offset the distortions in location and ownership choices induced by the revenue and labor taxes. The final term targets the national security externality directly and generates the expressions in Propositions 2 and 3.

A.3 Comparative statics

Figure 4: Optimal revenue, labor, and ownership wedges as a function of θ.

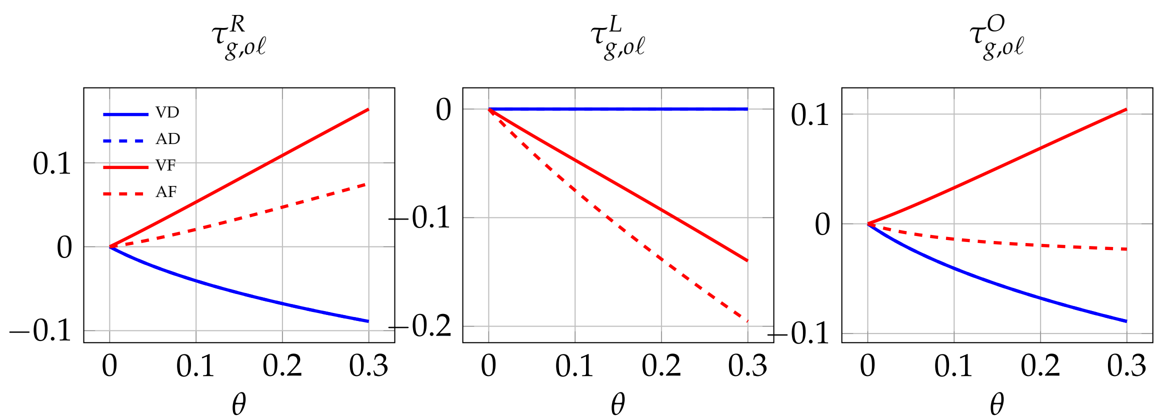

Figure 5: Cross-sectional appendix to Figure 2 for $\alpha_g = 0.75$.

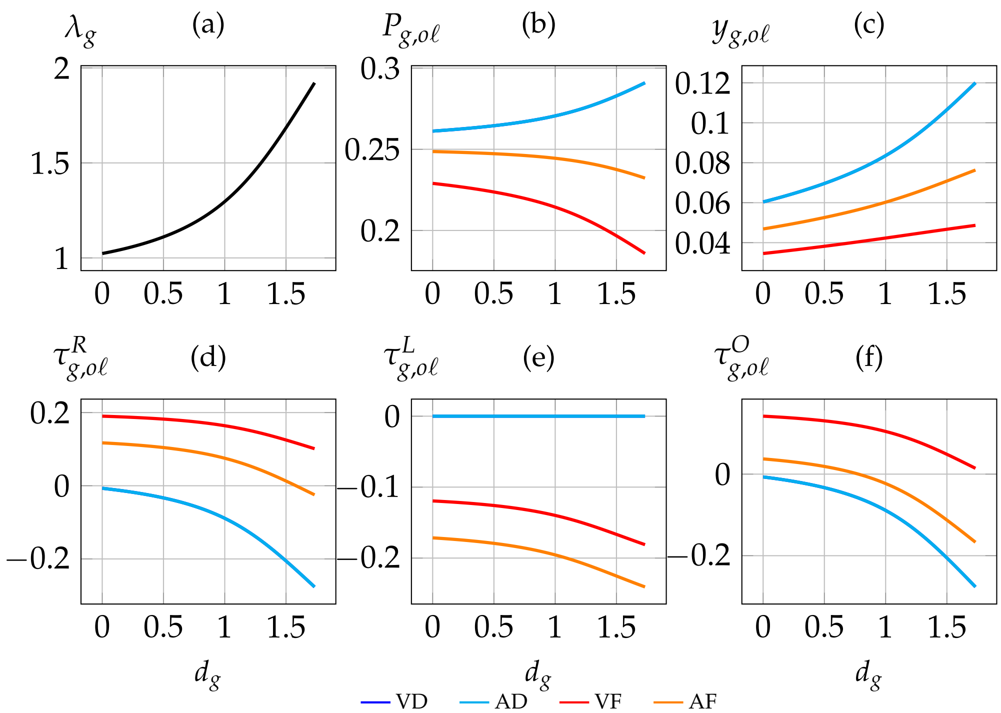

Note: d_g denotes peacetime government demand for good g. On the horizontal axis, d_g is reported relative to its baseline value, normalized so that the baseline equals 1. The figure reports optimal-policy outcomes across goods for the high-noncontractible calibration $\alpha_g = 0.75$, with conflict demand given by $d_g^C = 2d_g$.

Figure 6: Comparative statics across non-contractibility.

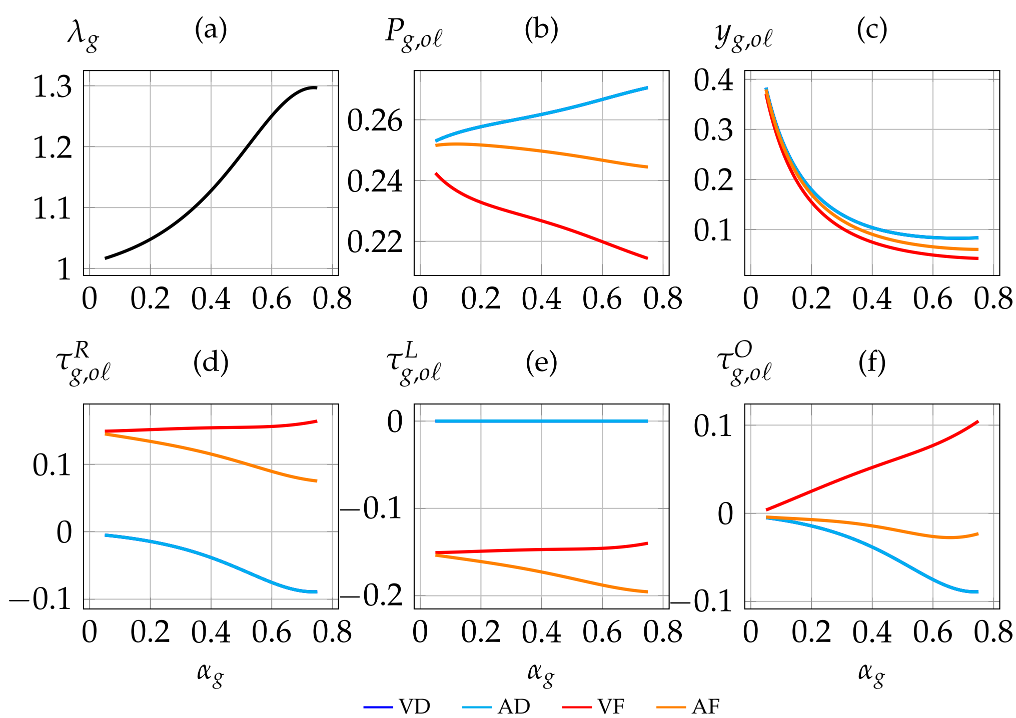

Note: This figure reports alpha comparative statics holding peacetime government demand fixed at its baseline value, $d_g = 0.025$. Conflict demand satisfies $d_g^C = 2d_g$. The panels report the conflict price ratio, organizational-mode shares, firm-level output, and the three wedges as α_g varies.

B Empirical Appendix

Figure 7: CFIUS filings over time

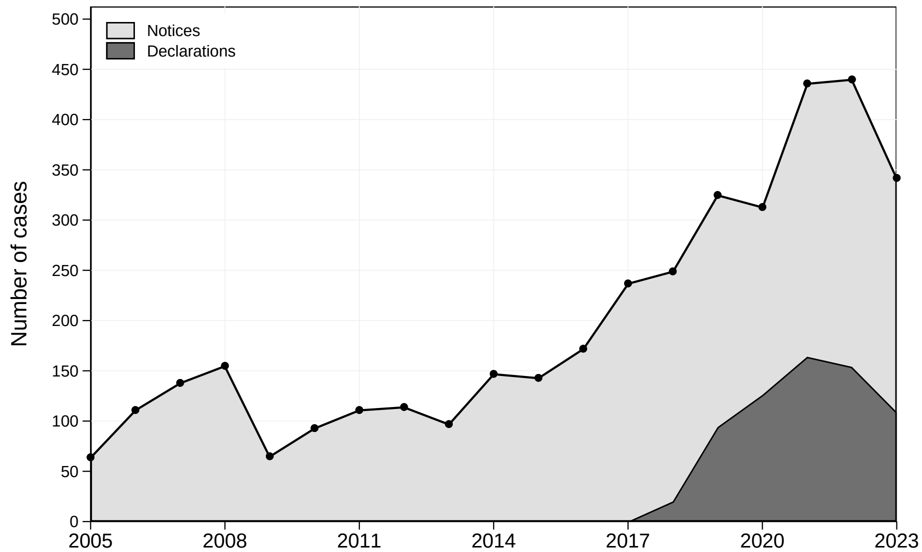

Notes: The figure displays the total number of covered transaction notices and declarations filed with CFIUS by calendar year. Declarations were introduced under the FIRRMA pilot program beginning in November 2018; the full declaration pathway took effect in February 2020. Source: CFIUS Annual Reports to Congress.

Figure 8: Share of CFIUS notices concluded with mitigation

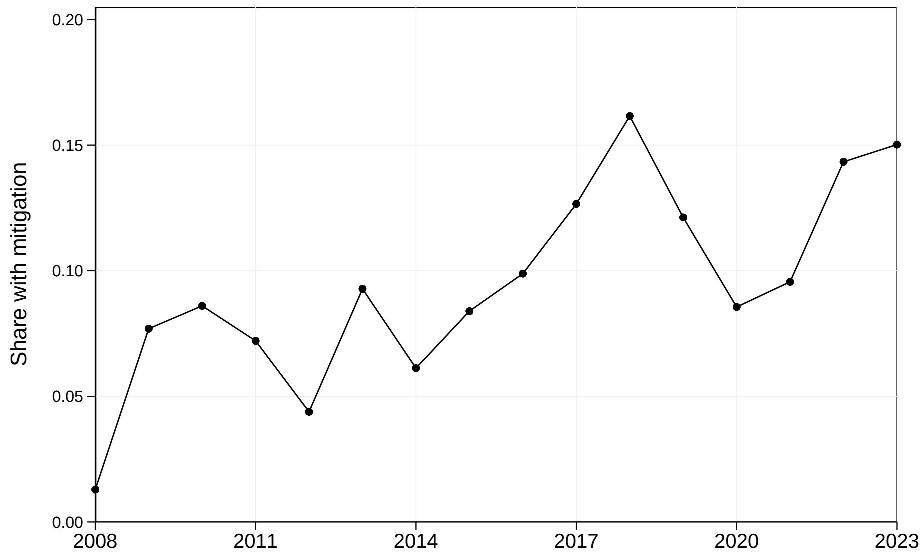

Notes: The figure displays the share of covered transaction notices for which CFIUS negotiated mitigation agreements (e.g., national security agreements or letters of assurance) as a condition for clearance, by calendar year. Source: CFIUS Annual Reports to Congress.

Table 2: Top Ten Manufacturing Industries by CFIUS Filings

NAICS	Industry	Notices	Decl.	Total
3344	Semiconductor and Other Electronic Component Mfg.	43	21	64
3364	Aerospace Product and Parts Mfg.	28	22	50
3345	Navigational, Measuring, Electromedical, and Control Instruments Mfg.	30	19	49
3342	Communications Equipment Mfg.	19	10	29
3329	Other Fabricated Metal Product Mfg.	11	12	23
3359	Other Electrical Equipment and Component Mfg.	14	3	17
3333	Commercial and Service Industry Machinery Mfg.	9	8	17
3341	Computer and Peripheral Equipment Mfg.	8	7	15
3254	Pharmaceutical and Medicine Mfg.	9	5	14
3332	Industrial Machinery Mfg.	9	5	14

Notes: Top ten 4-digit NAICS industries within manufacturing (NAICS 31–33) in terms of number of filings with CFIUS, 2021–2024. Source: CFIUS Annual Reports to Congress.

Figure 9: CFIUS filings by sector

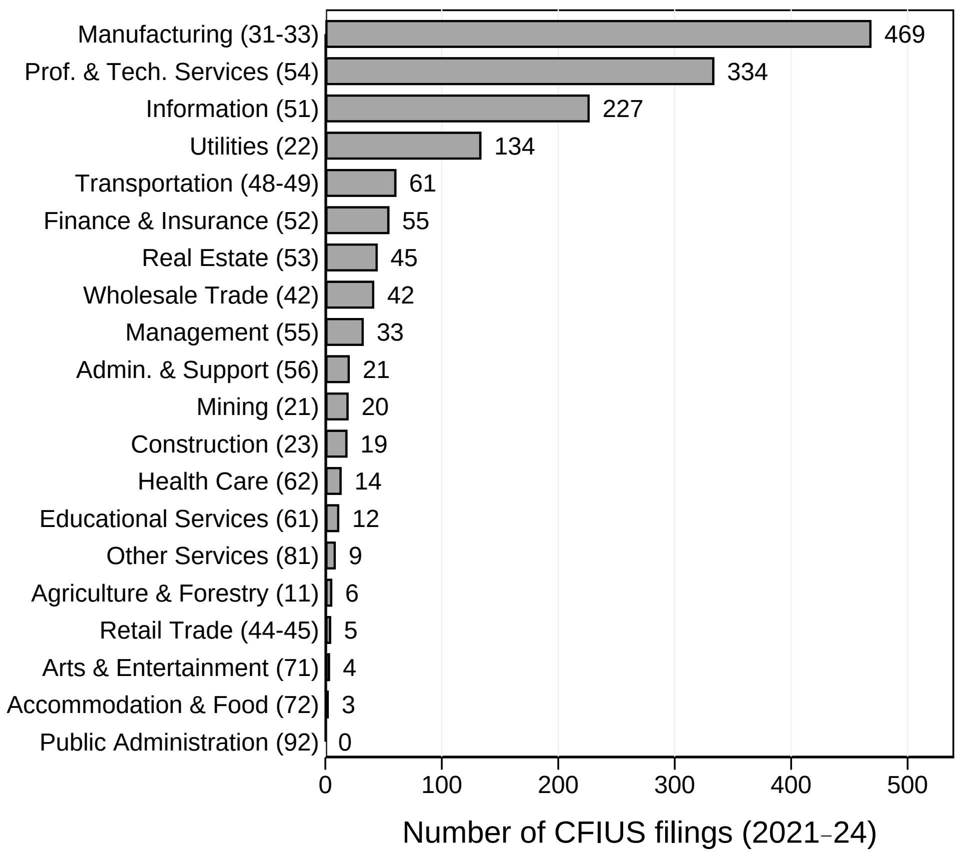

Notes: The figure displays the total number of CFIUS filings (notices plus declarations, 2021–2024) by two-digit NAICS sector, covering all NAICS 2017 sectors. Source: CFIUS Annual Reports to Congress.

Table 3: CFIUS Regulatory Intensity and Sector Characteristics

Dep. var.:	All filings	
	Baseline	+ Controls
	(1)	(2)
Defense demand, d_g	0.505	0.453
	(0.090)	(0.124)
Non-contract. input int., α_g	0.275	0.346
	(0.076)	(0.091)
$d_g \times \alpha_g$	0.265	0.349
	(0.079)	(0.110)
Capital intensity		0.216
		(0.111)
$d_g \times$ Capital int.		0.281
		(0.148)
Levchenko contract int.		0.055
		(0.078)
$d_g \times$ Levchenko int.		0.110
		(0.077)
Log total output		0.078
		(0.084)
M&A deal count (2021–24)		0.102
		(0.060)
R^2	0.529	0.638
Observations	85	85

Notes: Each column reports OLS estimates of (18) across four-digit NAICS industries in manufacturing. The dependent variable is the standardized number of CFIUS filings. Columns (1) and (2) correspond to columns (1) and (6) of Table 1, respectively. Column (2) additionally controls for log total output, the number of inbound cross-border M&A deals targeting U.S. firms in 2021–2024, and two alternative measures of non-contractible input intensity interacted with defense demand: the capital-labor ratio following Antràs (2003), and the negative of the Herfindahl-Hirschman Index of input spending shares following Levchenko (2007). Robust standard errors in parentheses.