Distribution Agreement

Papers on the Political Economy of the European Union

By

Joshua C. Fjelstul
Doctor of Philosophy

Political Science

Clifford J. Carrubba, Ph.D.
Advisor

Matthew J. Gabel, Ph.D.
Committee Member

Adam N. Glynn, Ph.D.
Committee Member

B. Pablo Montagnes, Ph.D.
Committee Member

Accepted:

Lisa A. Tedesco, Ph.D.
Dean of the James T. Laney School of Graduate Studies

Date

Papers on the Political Economy of the European Union

By

Joshua C. Fjelstul
M.A., Emory University, 2016
B.B.A., McCombs School of Business, The University of Texas at Austin, 2013
B.A., The University of Texas at Austin, 2013

Advisor: Clifford Carrubba, Ph.D.

An abstract of
A dissertation submitted to the Faculty of the
James T. Laney School of Graduate Studies of Emory University
in partial fulfillment of the requirements for the degree of
Doctor of Philosophy
in Political Science
2019

Abstract

Papers on the Political Economy of the European Union
By Joshua C. Fjelstul

This dissertation consists of three papers on the political economy of the European Union (EU). The central theme is how political institutions shape noncompliance.

In the first paper, I study how the politics of adjudication affect the implementation of policy. Under what conditions do government agencies more closely implement laws, and when do courts deter incorrect implementation? I model the strategic interaction between an implementing actor, a plaintiff, and a court over the implementation of a law. Implementing actors make concessions by implementing more complaint policies in order to avoid litigation. Courts are most effective at deterring noncompliance for intermediate levels of preference divergence between policy-makers and implementing actors. I test the observable implication of the model in the context of the EU using a novel dataset of implementation opportunities and noncompliance cases.

In the second paper, I study noncompliance in common markets. States create common markets to accrue consumer welfare gains. Given incentives to cheat to protect domestic firms from foreign competition, they create international regulatory regimes to manage noncompliance. I develop a formal model that explains how the politics of compliance in regulatory regimes systematically distorts the welfare gains that states accrue from developing common markets. The model predicts that regulatory regimes are most effective at enforcing compliance in sectors with intermediate levels of firm homogeneity in terms of productivity. The model also predicts the downstream consequences for the performance of individual firms and consumer welfare.

In the third paper, I study the how individuals evaluate complex economic policies. To what extent are aggregate preferences over complex economic policies consistent with individual rationality? This question has far-reaching implications for the coherence of economic policy. In this paper, I theorize how a rational individual would evaluate the following policy: Commission monitoring of member state compliance with the Stability and Growth Pact (SGP), which governs the European Economic and Monetary Union (EMU). I present empirical evidence from the European Sovereign Debt Crisis that, in a high-information environment, individuals express preferences over Commission monitoring that, in the aggregate, are consistent with individual rationality.

Papers on the Political Economy of the European Union

By

Joshua C. Fjelstul
M.A., Emory University, 2016
B.B.A., McCombs School of Business, The University of Texas at Austin, 2013
B.A., The University of Texas at Austin, 2013

Advisor: Clifford Carrubba, Ph.D.

A dissertation submitted to the Faculty of the
James T. Laney School of Graduate Studies of Emory University
in partial fulfillment of the requirements for the degree of
Doctor of Philosophy
in Political Science
2019

Acknowledgements

I'm very grateful to my committee — Cliff Carrubba, Matt Gabel, Adam Glynn, and Pablo Montagnes — for all of their comments, feedback, and advise on this project and others. I'm especially grateful to Cliff for the enormous amount of time he's spent with me since I arrived at Emory and for the confidence he's shown in my work. This project is so much better because of the huge investment he's made in me.

I'd like to thank Jim Caporaso, Greg Martin, Danielle Jung, Joe Jupille, Daniel Naurin, Shawn Ramirez, Eric Reinhardt, Dan Reiter, Jeff Staton, and Georg Vanberg either for their comments on various aspects of this project or for helping me to develop all of the skills that I have used in it.

I'd also like to thank Abbey for all of her unwavering support. (Skye and Iona too.)

This dissertation is dedicated to my parents. Without their unconditional love and support, this would never have been possible.

Contents

List of Tables

List of Figures

Paper 1:

Policy Implementation, Noncompliance, and the Judicial Impact of Courts

Abstract

Policy-makers depend on government agencies to implement laws, and courts adjudicate disputes over incorrect implementation. This paper analyzes how the politics of adjudication affect the implementation of policy. Under what conditions do government agencies more closely implement laws, and when do courts deter incorrect implementation? I model the strategic interaction between an implementing actor, a plaintiff, and a court over the implementation of a law. In equilibrium, implementing actors make concessions by implementing more complaint policies in order to avoid litigation. I find that courts are most effective at deterring noncompliance for intermediate levels of preference divergence between policy-makers and implementing actors. I test the observable implication of the model in the context of the European Union (EU) using a novel dataset of implementation opportunities and noncompliance cases.

1.1 Introduction

Policy-makers depend on other actors, like government agencies, to implement laws. However, the preferences of implementing actors can run counter to those of policy-makers. This creates an incentive to implement noncompliant policies, in which case policy-makers rely on high courts to correct their noncompliance. These courts, in turn, depend on other political actors, like the public and government agencies, to enforce their decisions (Carrubba 2005; Vanberg 2005; Carrubba 2009). Consequently, they are cautious in issuing adverse rulings: they anticipate noncompliance with their rulings by defendants, and are less likely to rule against defendants when the likelihood of noncompliance is high (Carrubba and Zorn 2010; Carrubba and Gabel 2015; Vanberg 2015).

This paper analyzes how the politics of adjudication affect the implementation of policy. Under what conditions do government agencies more closely implement laws, and when do courts deter incorrect implementation? To answer these questions, I present a formal model that predicts the conditions under which high courts will be most effective at deterring noncompliance by implementing actors. The model applies to constitutional courts, like the United States Supreme Court or the Bundesverfassungsgericht (the German Federal Constitutional Court), and quasi-constitutional international courts, like the Court of Justice of the European Union (CJEU), that have appellate jurisdiction over cases involving the incorrect implementation of laws.

I show that the impact of courts on policy outcomes depends not only on whether they are willing to rule against defendants, but also whether plaintiffs are willing to bring cases, and whether implementing actors adjust their behavior in anticipation of litigation. In equilibrium, courts deter noncompliance; implementing actors make policy concessions to avoid going to court. However, they do not come fully into compliance. Thus, neither implementing actors nor policy-makers achieve their first-best policy outcome. Courts are most effective at deterring noncompliance for intermediate levels of preference divergence between policy-makers and implementing actors.

This paper has important implications for recent research, particularly work on the effectiveness of quasi-constitutional international courts. A number of studies suggest that international courts are sensitive to the costs of compliance (Carrubba 2005, 2009; Gilligan and Johns 2012; Carrubba and Gabel 2015; Fjelstul and Carrubba 2018). Studies that do not account for this strategic behavior (e.g., Johns 2012; König and Mäder 2014; Johns 2015) will misidentify the conditions under which plaintiffs bring cases, and therefore the conditions under which courts deter noncompliance by implementing actors.

I test the key observable implication of the model in the context of the European Union (EU), where member states are responsible for implementing EU directives (i.e., secondary laws). The European Commission is responsible for monitoring the implementation of directives and can bring cases against member states. I focus my analysis on directives that harmonize industry regulations across the European Single Market. I develop a measure of preference divergence that is specific to harmonizing directives, create a sample of over 14,000 implementation opportunities (2004–2015), and code whether the Commission brings a noncompliance case *vis-à-vis* each implementation opportunity. I find evidence consistent with the model and inconsistent with an alternative that does not take into account the sensitivity of courts to noncompliance with their rulings.

1.2 The Strategic Behavior of Courts

Recent studies emphasize that constitutional courts and quasi-constitutional international courts are sensitive to whether defendants are likely to comply with adverse rulings (*ex post* compliance) (e.g., Staton and Moore 2011; Carrubba and Gabel 2017). This paper contributes to the literature on the judicial impact of these courts by providing a general formal theory of how they affect the implementation of policy that incorporates this insight into the strategic behavior of courts.

The literature on international courts argues that member states of international organizations create courts to solve collective action problems: states often benefit from cooperation but have incentives to unilaterally defect to avoid the costs of compliance. International courts serve as a fire alarm (i.e., a forum in which violations can be alleged) and an information clearing house (i.e., a forum in which interested actors can express their preferences). Since international agreements are incomplete contracts, and do not fully specify obligations under every possible set of circumstances, courts can improve cooperation by defining what constitutes an impermissible violation of the agreement (Carrubba and Gabel 2017). Courts enforce agreements by ruling against noncompliant member states when cooperation is mutually beneficial, but they permit noncompliance when the costs of compliance are so high that a defendant is unlikely to comply with an adverse ruling. As the costs of compliance increase, the probably of *ex post* noncompliance — the joint probability, conditional on a case, that a court rules against a defendant and the defendant complies with the adverse ruling — will decrease (Carrubba and Gabel 2015).

The literature on constitutional courts provides empirical evidence that is consistent with the argument that courts are sensitive to noncompliance with adverse rulings. Scholars have found that noncompliance depends on whether governments believe that the public can observe noncompliant behavior (Vanberg 2005). Courts use tools like press releases to increase public awareness of rulings to increase the costs of *ex post* noncompliance (Staton 2010). Courts also use citations to embed politically controversial rulings in established case law (Lupu and Fowler 2013; Lupu and Voeten 2012; Larsson et al. 2017).

Looking at the Court of Justice of the European Union (CJEU), specifically, which is the focus of my empirical test, recent studies have found evidence that the CJEU is sensitive to the preferences of member states (Carrubba and Gabel 2015; Larsson and Naurin 2016; Martinsen 2015; Blauberger and Schmidt 2017). Member states can use two mechanisms

— legislative override (Larsson and Naurin 2016; Martinsen 2015) and noncompliance (Carrubba and Gabel 2015) — to constrain the Court. The Court is more likely to rule against defendants in noncompliance cases when third-party member states support the plaintiff, as the defendant is more likely to comply *ex post* when other member states have an interest in enforcing an adverse ruling. The Court is less likely to rule against member state defendants in situations where institutional rules make legislative override easier (Larsson and Naurin 2016; Martinsen 2015) and when it is easier for member states to revise the EU Treaties (Castro-Montero et al. 2018).

The key takeaway from this literature is that courts do not make decisions in a vacuum. They recognize that their decisions only matter if defendants are willing to comply with them and work to anticipate when this is the case (Vanberg 2015; Carrubba and Gabel 2017; Blauberger and Schmidt 2017).[1] In this paper, I show that failing to account for the sensitivity of a court to *ex post* noncompliance (e.g., Johns 2012; König and Mäder 2014; Johns 2015) can yield substantively incorrect predictions about how international courts affect policy implementation.

1.3 Formal Model

To identify the conditions under which courts deter the incorrect implementation of laws *ex ante*, I model the strategic interaction between an implementing actor, a plaintiff, and a high court over the implementation of a law. Consistent with the recent literature on international courts, I assume that the court is sensitive to the probability of *ex post* compliance. I show that the plaintiff anticipates the likelihood that the court will rule against a defendant and select out of cases when the likelihood of *ex post* compliance is low.[2] Thus, if

[1] See Vanberg (2015), Carrubba and Gabel (2017), and Blauberger and Schmidt (2017) for a more in-depth review of recent literature.

[2] In the context of the EU, scholars have suspected that the Commission strategically chooses which noncompliance cases to pursue based on the likelihood of success (Mbaye 2001; Börzel 2003; Thomson, Torenvlied and Arregui 2007; Hartlapp and Falkner 2009; Steunenberg and Rhinard 2010). Recent empirical work finds

courts are sensitive to noncompliance with their rulings, that leads to systematic bias in the noncompliance cases that get litigated.

Implementing actors anticipate this strategic behavior by plaintiffs and make concessions to policy-makers by moderating their noncompliance; they make larger concessions when the likelihood of litigation is higher. Thus, even if a court does not rule against a noncompliant defendant in a particular case, because the court believes the member state is unlikely to comply, the fact that it could have prompts the implementing actor to implement a more compliant policy in the first place. This concession, which represents the judicial impact of the court on policy outcomes, is largest for intermediate levels of preference divergence between policy-makers and implementing actors.

The model applies to a variety of institutional settings. The high court could be a domestic constitutional court or a quasi-constitutional international court (i.e., a court of last resort) that has appellate jurisdiction over cases involving the incorrect implementation of laws by implementing actors. The policy-maker could be a legislative body, an executive body with delegated policy-making authority (e.g., an executive agency or the European Commission), or an international decision-making body (e.g., the Council of the European Union). In unitary states, the implementing actor could be an executive agency; in federal states, it could be a regional government; in supranational organizations (e.g., the European Union), it could be a member state government. The plaintiff could be a private actor (e.g., a citizen, a firm, a non-governmental organization, etc.) or an institution that is tasked with monitoring and enforcing compliance (e.g., the European Commission).

1.3.1 Order of Play

The model has three players: an implementing actor, a plaintiff, and a high court. I model the court as a reduced-form actor. There is also a non-strategic policy-maker that sets a

evidence that the Commission does, in fact, strategically choose cases (König and Mäder 2014; Fjelstul and Carrubba 2018).

de jure policy, $p_m \in \mathbb{R}^1$, that the implementing actor must implement. The policy set by the policy-maker p_m is exogenous. The implementing actor chooses a *de facto* policy that implements the *de jure* policy. There are two versions of the *de facto* policy: the *ex ante* policy, which is the one implemented prior to any litigation, and the *ex post* policy, which is the one implemented after any litigation. The *ex ante* policy is $p_0 \in \mathbb{R}^1$ and the *ex post* policy is $p_1 \in \mathbb{R}^1$. A correctly implemented *ex ante* policy is $p_0 = p_m$.

Figure 1.1 summarizes the order of play. The game starts with the implementing actor choosing an *ex ante* policy. Any *ex ante* policy that does not exactly equal the *de jure* policy, $p_0 \neq p_m$, is noncompliant. The larger the absolute difference, the more noncompliant the policy; thus, in equilibrium, *ex ante* noncompliance is $\left| p_0^* - p_m \right|$. The plaintiff observes this policy and decides whether or not to bring a noncompliance case.[3]

Definition 1. *Ex ante* noncompliance is the degree to which the *ex ante* policy chosen by the implementing actor differs from the *de jure* policy, $\left| p_0^* - p_m \right|$.

I model the court as a reduced-form player. Consistent with the recent literature on the strategic behavior of constitutional courts and quasi-constitutional international courts (e.g., Vanberg 2001, 2005; Carrubba 2005, 2009; Gilligan and Johns 2012; Vanberg 2015; Carrubba and Gabel 2015; Fjelstul and Carrubba 2018), the court has preferences over *ex post* compliance with adverse rulings; it is hesitant to rule against a defendant that is unlikely to comply with the ruling.

If the plaintiff brings a case, there is a conditional probability of *ex post* compliance, $h(c)$, which is the joint probability that the court rules against the implementing actor and that the implementing actor complies with that adverse ruling. This joint probability depends on the cost of compliance, $c > 0$, which is the difference in utility (for the implementing actor) between fully complying (i.e., correctly implementing the *de jure* policy) and implementing an optimal policy in a counterfactual model without a court to adjudicate disputes over

[3] Similar to Fjelstul and Carrubba (2018), I do not allow the plaintiff to infer a violation where one does not exit; the plaintiff and the court both have perfect information about the degree to which the *ex ante* policy is noncompliant.

1. **Implementation**: The implementing actor choses an *ex ante* policy, p_0, that implements the *de jure* policy, p_m, set by a policy-maker.
2. **Litigation**: The plaintiff brings a noncompliance case against the implementing actor at cost k or not.
3. **Adjudication**: If the plaintiff brings a case, there is *ex post* compliance with probability $h(c)$, in which case the *ex post* policy is $p_1 = p_m$, and *ex post* noncompliance with probability $1 - h(c)$, in which case the *ex post* policy is $p_1 = p_0$.

Figure 1.1. Order of Play

Note: This figure shows the order of play. Note that the probability of *ex post* compliance, $h(c)$, is the joint probability that the court rules against the implementing actor and that the implementing actor complies with the ruling. The probability of *ex post* noncompliance, $1 - h(c)$, is the probability that the court rules in favor of the implementing actor plus the joint probability that the court rules against the implementing actor and that the implementing actor does not comply with the ruling.

implementation. Consistent with the literature (e.g., Carrubba 2005, 2009; Gilligan and Johns 2012; Carrubba and Gabel 2015), as the cost of compliance increases, the probability of *ex post* compliance decreases, $h'(c) < 0$. The implementing actor becomes less likely to comply with an adverse ruing and, anticipating this, the court becomes less likely to rule against the defendant.[4]

Definition 2. Conditional on a noncompliance case, the probability of *ex post* compliance, $h(c)$, is the joint probability that the court rules in favor of the plaintiff and that the implementing actor complies with that adverse ruling.

Assumption 1. The probability of *ex post* compliance, conditional on a noncompliance case, is decreasing in the cost of compliance, $h'(c) < 0$.

If there is *ex post* compliance, the implementing actor's *ex post* policy is the *de jure* policy, $p_1 = p_m$, regardless of whatever its *ex ante* policy was. If there is not *ex post* compliance, then the implementing actor's *ex post* policy is the same as its *ex ante* policy, $p_1 = p_0$, just like when the plaintiff does not bring a case. Thus, the *ex post* policy can only take one

[4] Fjelstul and Carrubba (2018) model a strategic court in the same way.

of two values, $p_1 \in \{p_0, p_m\}$.[5] In equilibrium, *ex post* compliance is the absolute difference between the *ex post* policy chosen by the implementing actor and the *de jure* policy set by the policy-maker, $\left| p_1^* - p_m \right|$. If the plaintiff does not bring a case, the *ex post* policy is the same as the *ex ante* policy and *ex ante* noncompliance is the same as *ex post* noncompliance. Since litigation is a stochastic process, *ex post* noncompliance is an expectation, $\left| E[p_1^*] - p_m \right|$.

Definition 3. Expected *ex post* noncompliance is the degree to which the *ex post* policy produced by the litigaiton process differs from the *de jure* policy set by the policy-maker in expectation, $\left| E[p_1^*] - p_m \right|$.

1.3.2 Preferences

The implementing actor has a preference over its *ex post* policy p_1. Its utility is given by a quadratic loss function, defined over p_1: $u_A(p_1) = -(p_1 - p_i)^2$, where p_i is its ideal point. I assume, without loss of generality, that the implementing actor prefers a policy to the right of the *de jure* policy, $p_i > p_m$, which I normalize, $p_m = 0$. Preference divergence, in the context of the model, is the absolute difference between the implementing actor's ideal point and the *de jure* policy set by the policy-maker, $\left| p_i - p_m \right|$. Preference divergence will be the primary exogenous parameter of interest.

Definition 4. Preference divergence is the absolute difference between the ideal point of the implementing actor and the *de jure* policy, $\left| p_i - p_m \right|$.

The plaintiff also has a preference over the *ex post* policy. If the plaintiff brings a case, it pays a cost of litigating, $k > 0$, which is private information. The plaintiff learns k before deciding whether or not to bring a case. This cost is drawn from a probability density

[5] I assume that the implementing actor either fully complies with an adverse ruling, $p_1 = p_m$, or fully ignores it and keeps whatever policy it has chosen *ex ante*, $p_1 = p_0$. Allowing the implementing actor to partially comply *ex post* would not substantively change its *ex ante* policy-implementation behavior in equilibrium. Even if the implementing actor only partially complies *ex post*, that outcome will still incentivize the implementing actor to try to avoid litigation by making a policy concession. This is true as long as the concession that the implementing actor makes *ex post* is worse than the concession it makes *ex ante*.

function $f(k)$ with a finite mean and variance. The support of $f(k)$ is $k \in [0, \infty)$. It has a cumulative distribution function $F(k)$, where $F'(k) > 0$. Low-cost litigation is more likely than high-cost litigation, which implies that the probability density function $f(k)$ is decreasing in k, $F''(k) = f'(k) < 0$. I also assume that the plaintiff prefers the *de jure* policy (i.e., that the preferences of the policy-maker and the plaintiff are aligned). Thus, the plaintiff's utility is as follows:

$$u_P(p_1) = \begin{cases} -(p_1 - p_m)^2 & \text{if there is no case} \\ -(p_1 - p_m)^2 - k & \text{if there is a case and no } \textit{ex post} \text{ compliance} \\ -k & \text{if there is a case and } \textit{ex post} \text{ compliance.} \end{cases}$$

To summarize the information structure, the model assumes (a) that noncompliance is always strategic, and not accidental; (b) that the plaintiff and the court are fully informed about the degree to which the implementing actor has correctly implemented the *de jure* policy set by the policy-maker; and (c) that the costs of litigation are private information, but are known to the plaintiff before the plaintiff decides whether or not to bring a case.[6] Thus, if the court sides with the implementing actor, it is making a political decision to issue a ruling contrary to the legal merits of the case.

1.3.3 Equilibrium Behavior

I derive a unique subgame perfect equilibrium by backwards induction. See Appendix A for formal proofs. I present the equilibrium in reverse-chronological order. Note that all play is on the equilibrium path.

I start by considering the counterfactual: the policy the implementing actor would choose if there were no court to adjudicate disputes. In the counterfactual, there are no institutional

[6] Fjelstul and Carrubba (2018) similarly assume that the plaintiff and court know whether noncompliance has occurred and that the costs of litigation are private information, but in that model, the plaintiff is uncertain whether compliance is accidental or strategic.

constraints on the implementing actor (i.e., no prospect of costly litigation). Thus, the implementing actor simply chooses to implement its preferred policy. Since there is no prospect of litigation, the *ex ante* and *ex post* policy are one in the same.

Lemma 1. In the counterfactual, the implementing actor is free from institutional constraints and chooses an optimal policy equal to its ideal point, $p^* = p_i$.

The behavior of the court depends on the cost of compliance, which is an equilibrium quantity. The cost of compliance is how much worse off the implementing actor would be if it fully complied instead of choosing the policy it would choose in the counterfactual. This is the absolute difference in utility between complying with the *de jure* policy and choosing an optimal policy in the counterfactual.

Lemma 2. The implementing actor's cost of compliance in equilibrium is the absolute difference in its utility between correctly implementing the *de jure* policy set by the policy-maker, $p_1 = p_m$, and choosing an optimal policy in the counterfactual, $p^* = p_i$: $c^* = \left| \, u_A(p_m) - u_A(p_i) \, \right| = p_i^2$.

The plaintiff observes *ex ante* noncompliance and decides whether or not to bring a noncompliance case. Intuitively, the plaintiff brings a case when the cost of litigation is sufficiently small relative to the degree of *ex ante* noncompliance and the conditional probability of *ex post* noncompliance.[7]

Proposition 1. The plaintiff observes *ex ante* noncompliance, $\left| \, p_0^* - p_m \, \right|$, anticipates the conditional probability of *ex post* noncompliance, $h(c^*)$, and brings a case when the cost of litigation is sufficiently low, $k < k^*$, where $k^* = (p_0^*)^2 h(p_i^2)$. The probability that the plaintiff brings a case is $\Pr(k < k^*) = F(k^*)$.

The implementing actor anticipates the probability that the plaintiff will bring a case and the conditional probability of *ex post* compliance. Recall that the implementing actor has

---

[7] This is consistent with the recent literature on the European Commission (Konig and Mader 2014; Fjelstul and Carrubba 2018).

a preference over the *ex post* policy. Since litigation is a stochastic process, in equilibrium, the implementing actor chooses an *ex ante* policy that maximizes its utility in expectation. The cost of bringing a case for the plaintiff is private information, so the implementing actor only knows the probability of a case.

If the plaintiff brings a case, with probability $F(k^*)$, there is some probability that the court will rule against the implementing actor and that the implementing actor will comply, $h(c^*)$. In this case, the implementing actor implements the *de jure* policy and receives $u_A(p_m)$. With probability $1 - h(c^*)$, there is *ex post* noncompliance.[8] In this case, the implementing actor's *ex post* policy is equal to its *ex ante* policy, $p_1 = p_0$, and the implementing actor receives $u_A(p_0)$. If the plaintiff does not bring a case, with probability $1 - F(k^*)$, the implementing actor's *ex post* policy is also equal to its *ex ante* policy. Thus, the implementing actor's expected utility is as follows:

$$E\big[u_A(p_0)\big] = F(k^*)\big[h(c^*)u_A(p_m) + (1 - h(c^*))u_A(p_0)\big] + (1 - F(k^*))u_A(p_0). \tag{1.1}$$

The implementing actor's expected utility captures a basic tradeoff between improving the policy (i.e., moving the policy closer to its ideal point) and triggering a court battle, which could end in *ex post* compliance (the worst-case policy outcome for the implementing actor). The implementing actor can reduce the likelihood of litigation by compromising — by implementing a policy that is closer to the *de jure* policy — but this comes at the cost of a worse policy outcome if there is not a case (or if there is and the court rules in favor of the implementing actor).[9] In equilibrium, the implementing actor chooses the *ex ante* policy that optimally balances this tradeoff (i.e., there point where the marginal increase in its utility equals the marginal decrease).

[8] Recall that there are two ways to get *ex post* noncompliance, conditional on a case. Either the court does not rule against the implementing actor, or it does and the implementing actor ignores the ruling.

[9] This expected utility equation is a globally concave, single-peaked function, which guarantees a unique global maximum.

Proposition 2. The implementing actor chooses an optimal *ex ante* policy p_0^*. There is a unique solution, but there is not a closed-form expression. The implicit solution for p_0^* is given by the following first-order condition (FOC):

$$2p_0^2 h(p_i^2)^2(p_0 - 2p_i)F'(p_0^2 h(p_i^2)) + 2(p_0 - p_i)\Big[h(p_i^2)F(p_0^2 h(p_i^2)) - 1\Big] = 0.$$

In equilibrium, the implementing actor's optimal *ex ante* policy is a concession by the implementing actor to the policy-maker; it is a compromise between its ideal point and the *de jure* policy.[10] The effect of the court on the *ex ante* policy (and on *ex ante* compliance) in equilibrium is the size of this optimal concession, which depends on the level of preference divergence between the policy-maker and the implementing actor.

Proposition 3. The court incentivizes the implementing actor to make a policy concession, which is the absolute difference between the optimal *ex ante* policy and the optimal policy in the counterfactual, $\left| p_0^* - p^* \right|$. This is the judicial impact of the court on policy implementation.

This behavior is intuitive. The court constrains the policy-implementation behavior of the implementing actor. The prospect of losing in court incentivizes the implementing actor to trade a worse policy for a lower probability of going to court (i.e., a lower probability of enforced compliance).[11] In short, the shadow of enforced compliance induces more compliant behavior by the implementing actor. The question, though, is not only whether a court can improve the *ex ante* compliance (if the court matters at all, this should be true), but also the conditions under which the court is most effective at deterring *ex ante* noncompliance. In other words, when is the optimal concession larger or smaller?

[10] In other words, $p_m < p_0^* < p_i$.

[11] In model notation, the probability of enforced compliance is $F(k^*)h(c^*)$.

1.3.4 Comparative Statics

I use this equilibrium to analyze how preference divergence affects the behavior of the plaintiff and the implementing actor. The key prediction of the model is that the judicial impact of a court on policy outcomes will be greatest for intermediate levels of preference divergence between policy-makers and implementing actors. I show that if the court were not sensitive to noncompliance with adverse rulings (e.g., if going to court were a simple random lottery) we would reach an entirely different conclusion: the impact of the court in equilibrium would be strictly increasing in preference divergence. In institutional settings where the implementing actor is a member state of an international organization, this creates an incentive to exit the regime to avoid enforced compliance over the long term.

This key prediction is not directly testable because there is no plausible control group for the counterfactual (i.e., policy implementation without a court). However, the model also predicts that the probability of litigation is highest for intermediate levels of preference divergence, and this prediction is testable. In fact, this is the reason why the judicial impact of the court is greatest under these conditions. Conditional on empirical support for the model, if we assume that the model captures the data-generating process, we can also infer (a) the degree to which a court is sensitive to the cost of compliance, and (b) the severity of the compliance deficit caused by strategic noncompliance.

I present the comparative statics in reverse-chronological order. See Appendix A for proofs. I confirm all comparative statics using a computational simulation. Figure 1.2 summarizes all of the comparative statics using the output of this computational simulation. To sign the comparative statics, it will be necessary to choose functional forms for F and h. I assume that the functional form of F is the CDF of the exponential distribution and that the functional form of h is the logistic function.[12]

[12] I use Monte Carlo simulations to demonstrate the robustness of the results to variation in the rate parameter of the exponential distribution, the shape parameter of the logistic function, and the location parameter of the logistic function.

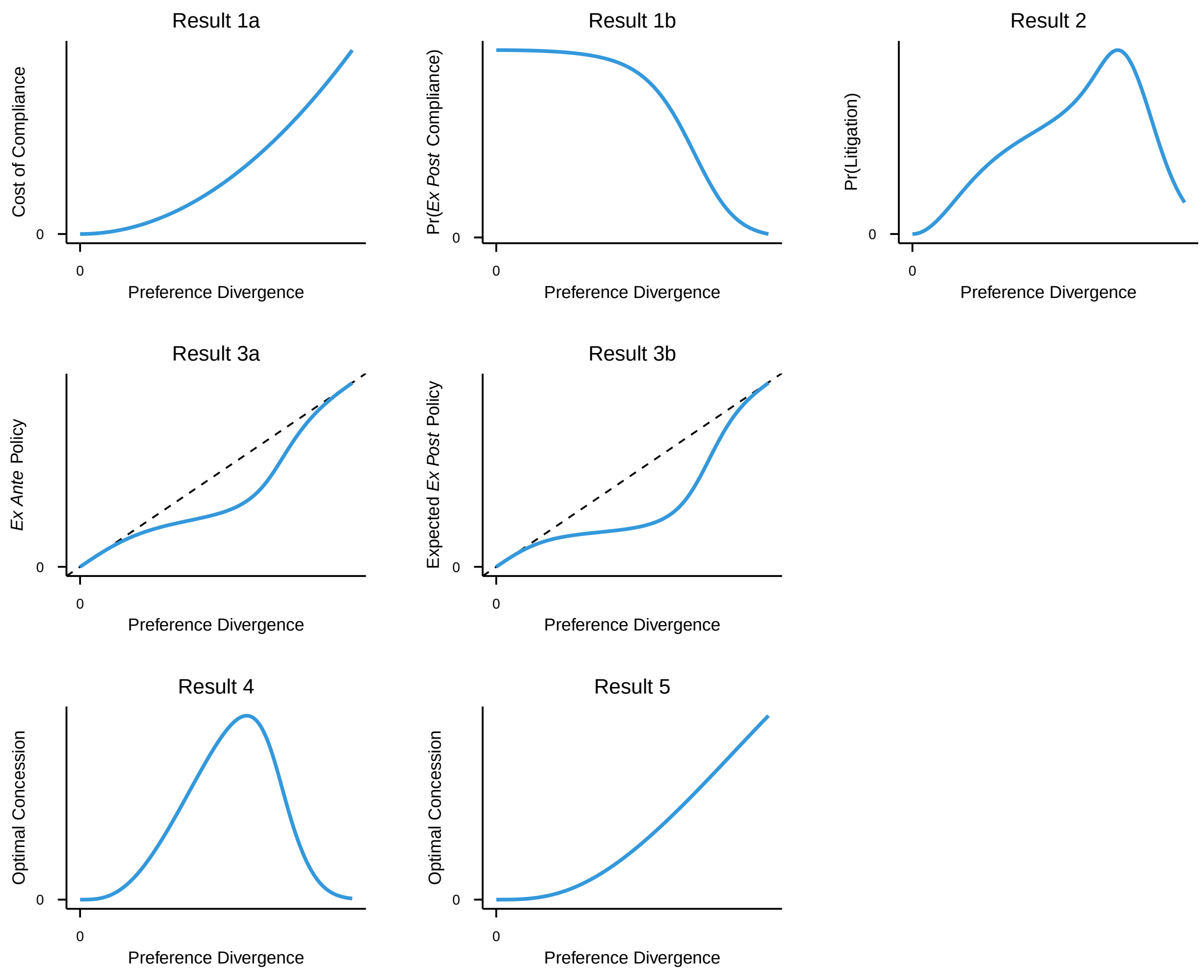

Figure 1.2. Summary of Comparative Statics

Note: This figure summarizes the comparative statics using a numerical simulation of the model. I use an exponential distribution for $F(k)$ and a logistic function for $h(c)$.

I start with the equilibrium behavior of the court: the probability of *ex post* compliance, conditional on a case, is decreasing in preference divergence. Intuitively, as preference divergence increases, the cost of compliance increases (Figure 1.2, Result 1a); when the implementing actor and the policy-maker disagree, it is more costly for the implementing actor to fully comply with the *de jure* policy. Consequently, the conditional probability of *ex post* compliance decreases (Figure 1.2, Result 1b), according to Assumption 1. The court becomes more hesitant to rule against the implementing actor out of concern that the implementing actor will ignore its ruling.

Result 1a. In equilibrium, as preference divergence increases, the implementing actor's cost of compliance is increasing.

Result 1b. In equilibrium, as preference divergence increases, the probability of *ex post* compliance, conditional on a case, is decreasing.

Turning next to the plaintiff, the probability of litigation is largest for intermediate levels of preference divergence (Figure 1.2, Result 2). When preference divergence is low, the plaintiff has little incentive to bring a noncompliance case. Noncompliance is minimal, so the value of correcting noncompliance is low relative to the costs of litigating. The potential benefit of bringing a case does not justify the costs of litigation. As preference divergence increases, the costs of compliance also increases, creating an incentive for the implementing actor to implement a less compliant *ex ante* policy (see Result 3a below). For the plaintiff, which prefers the *de jure* policy, the value of correcting noncompliance is increasing, so the costs of bringing a case are more likely to be justified. At the same time, however, the probability of *ex post* compliance, conditional on a case, is dropping (Assumption 1). As the cost of compliance increase, the court becomes less likely to rule against the implementing actor, which deters the plaintiff from bringing a case.

Thus, as preference divergence increases, the plaintiff faces a trade-off: the benefits of successful litigation are increasing, but the probability of successful litigation is decreasing. This trade-off produces a non-monotonic effect: the probability that the plaintiff brings a case is increasing as the benefits of correcting noncompliance increase, but past a certain point, the costs of litigating are no longer justified, given the falling probability of successful litigation, and the probability that the plaintiff brings a case starts to decline. The plaintiff is most likely to bring a case for intermediate levels of preference divergence. There is a systematic bias in the types of cases that get litigated.

Result 2. In equilibrium, as preference divergence increases, the probability that the plaintiff brings a noncompliance case is increasing, then decreasing.

Turning to the implementing actor, as preference divergence increases, the *ex ante* policy that the implementing actor chooses diverges from the *de jure* policy (Figure 1.2, Result 3a). As preference divergence increases, the cost of compliance increases; when the implementing actor and the policy-maker disagree, it is more costly for the implementing actor to fully comply with the *de jure* policy. This creates an incentive for the implementing actor to diverge from the *de jure* policy. This effect carries over to the *ex post* policy in expectation (Figure 1.2, Result 3b).

Result 3a. In equilibrium, as preference divergence increases, the optimal *ex ante* policy that the implementing actor chooses is increasing. Thus, the *ex ante* noncompliance deficit is increasing.

Result 3b. In equilibrium, as preference divergence increases, the *ex post* policy produced by the litigation process is increasing in expectation. Thus, the *ex post* noncompliance deficit is increasing in expectation.

The bias in which cases are litigated affects the degree to which the court can constrain the *ex ante* policy chosen by the implementing actor and the degree to which the court affects the *ex post* policy. Despite the fact that *ex ante* and *ex post* noncompliance are both increasing in preference divergence (Results 3a and 3b), the ability of the court to constrain the behavior of the implementing actor is not strictly decreasing in preference divergence. Instead, specifically because of this bias, the judicial impact of the court is larger for intermediate levels of preference divergence (Figure 1.2, Result 4).[13]

The intuition is that the implementing actor anticipates the probability of the plaintiff bringing a case, and when that probability is high, it is willing to make a larger concession to reduce the probability of litigation. Recall that the implementing actor wants to avoid litigation because going to court risks *ex post* compliance (i.e., the worst-case policy out-

[13] Looking at Panel 4 in Figure 1.2, the size of the optimal concession is the vertical difference between the solid line (indicating the optimal *ex ante* policy in equilibrium) and the dashed line (indicating the optimal policy in the counterfactual).

come). When preference divergence is minimal, the government wants to comply, and there is little tension. When preference divergence is substantial, on the other hand, the optimal concession disappears because the implementing actor knows the plaintiff is unlikely to bring a case. If the plaintiff did bring a case, the court, knowing the implementing actor would be unlikely to come into compliance with an adverse ruling, would likely rule in favor of the implementing actor, despite the merits of the case.

Result 4. In equilibrium, as preference divergence increases, the effectiveness of the court at deterring noncompliance (i.e., the size of the optimal concession) increases, then decreases.

This result is driven entirely by the systematic bias in which cases the plaintiff chooses to pursue (Result 2), which is driven by the fact that the court is sensitive to noncompliance with adverse rulings. If going to court were a simple random lottery, in which the court rules probabilistically and the implementing actor complies with adverse rulings probabilistically, the optimal concession would be strictly increasing in preference divergence (Figure 1.2, Result 5). In fact, it would be increasing at an increasing rate. This is true regardless of how likely the implementing actor is to comply with an adverse ruling.

We can model this by assuming that the probability of *ex post* compliance, conditional on the plaintiff brining a case, is a constant, $h(c^*) = w$, where w is the probability that the court rules in favor of the plaintiff. We can think of w as capturing preference divergence between the court and the policy-maker. When w approaches 1, the court is predisposed to agree with the policy-maker, and when w approaches 0, the court is predisposed to agree with the implementing actor. Regardless of the value of w, the optimal concession is always increasing in preference divergence.

Result 5. If the court were not sensitive to the costs of compliance, and the probability of *ex post* compliance, conditional on a case, were a constant, the judicial impact of the court on policy implementation (i.e., the size of the optimal concession) would be strictly increasing in preference divergence.

In sum, the model predicts that courts are most effective at deterring *ex ante* noncompliance for intermediate levels of preference divergence between implementing actors and policy-makers. If we do not take into account the political of adjudication (i.e., the sensitivity of a court to *ex post* noncompliance with adverse rulings), we get an entirely different result: the judicial impact of the court on policy implementation would be strictly increasing in preference divergence.

1.4 Empirical Analysis

I test the key empirical implication of the model — that the probability of litigation is increasing, then decreasing in preference divergence (Result 2) — in the context of the European Union (EU).[14] The EU is a multi-level supranational organization. The EU legislative institutions — the Council of the European Union (Council) and the European Parliament (EP) — create policy with the input of member states. The primary legislative instrument in the EU is a directive. Directives specify a policy outcome that all member states must achieve, but allow member states considerable flexibly in terms of how to achieve it. EU member states are responsible for implementing EU directives by transposing them into national law by the stated deadline.

The EU uses a centralized monitoring system to monitor and enforce compliance with EU law. The European Commission uses a multi-stage procedure called the infringement procedure to manage noncompliance. If the Commission suspects that an EU member state has committed an infringement, it can initiate an infringement case by sending a letter of formal notice to the member state. If the member state does not come into compliance, the Commission can send a reasoned opinion, which lays out a formal legal argument against the member state. If the member state still does not come into compliance, the Commission

[14] Since the counterfactual is not observable, we cannot test the effect of preference divergence on the size of the optimal concession. Result 2 is an observable implication because we can identify opportunities for implementing actors to implement a policy and we can observe whether plaintiffs bring cases *vis-à-vis* each implementation opportunity.

can refer the case to the Court of Justice of the European Union (Article 258 TFEU). The CJEU then determines whether the member state has committed an infringement. The CJEU issues a binary ruling on each count. If the CJEU rules in favor of the Commission, the member state must come into compliance.

Since Result 2 is a non-monotonic prediction, the sign of the average marginal effect that we should expect to recover by estimating a statistical model on data will depend on the region of the equilibrium space that the data come from, which is unknowable. In the context of the EU, however, there are theoretical reasons to expect that the data come from the region where the effect is negative, which is the region in which the model makes a prediction that differs from the prediction of an alternative model in which the court is not sensitive to *ex post* compliance.

In the model, the cost of litigation deters the plaintiff from bringing cases when preference divergence is low. As preference divergence increases, bringing a case is more likely to be worth the cost, resulting in a positive effect. If bringing a case were costless, the predicted effect would always be negative. The Commission handles thousands of noncompliance cases at the same time, which implies that the marginal cost of bringing a case is very low. Thus, we should expect to be in the region of the equilibrium space where the effect is negative. Finding a negative effect would be discriminating, but finding a positive effect would be consistent with both.

1.4.1 Research Design

To test Result 2, I create a sample of implementation opportunities — opportunities for EU member states to implement EU directives — and estimate the unconditional probability that the Commission brings an infringement case *vis-à-vis* each implementation opportunity.[15] We cannot observe noncompliance directly with the data that are available, but this is not

[15] A number of studies look at the likelihood of infringement cases, conditional on opportunities to transpose directives, but they generally look at small samples and do not consider the strategic behavior of the Commission (e.g., Thomson, Torenvlied and Arregui 2007).

necessary to test this prediction because the model predicts the unconditional probability of litigation, given an implementation opportunity.

I focus on EU directives that harmonize technical regulations. I develop a novel measure of preference divergence that is specific to this category of directives (see below). Technical regulations are economic regulations that restrict what kinds of goods can legally be sold.[16] Harmonization is the process of standardizing technical regulations across EU member states, which facilitates intra-EU trade by reducing the regulatory burden on firms and making it easier for firms to sell their products in other member states.[17] Over time, the EU has attempted to harmonize technical regulations in more and more industries.

I start by creating a novel dataset of the universe of implementation opportunities for the period 2004–2015.[18] I code one implementation opportunity per member state per directive. I do not code implementation opportunities for directives that member states implement as part of the accession process. In other words, I only code an implementation opportunity if a member state has already joined the EU by the date the directive was published. I collect metadata on 1,019 directives published between May 1, 2004 (the date that the EU officially expanded to 25 member states) and December 31, 2015 from EUR-Lex, the Commission's online database of EU legal documents. This is the universe of directives during this period. I exclude a small number of directives that are not addressed to all member states. There are a total of 26,684 implementation opportunities in the dataset.

I use regular expressions to extract the subject matter classifications that the Commission assigns to each directive. Since my measure of preference divergence applies specifically to

[16] According to the Commission, technical regulations include (a) technical specifications, (b) rules on providing services, and (c) regulations that prohibit the manufacture, importation, or marketing of a product. A technical specification is a required characteristic of a product, such as "dimension, labelling, packaging, [or] level of quality."

[17] According to the Commission, "Many products on the EU market are subject to harmonized rules that protect consumers, public health, and environment. Harmonized rules preclude the adoption of possibly divergent national rules and ensure the free circulation of products within the EU. Some sectors are still governed by national provisions however. The principle of free movement of goods ensures that these provisions do not lead to the creation of unjustified barriers to trade."

[18] This period is limited by the availability of case-level data on infringement cases.

harmonization directives, I only include directives that the Commission codes as dealing with harmonization (55.27 percent of all implementation opportunities).[19] The final dataset includes a total of 14,747 implementation opportunities.

For each implementation opportunity, I code whether the Commission opens an infringement case against the member state by sending a letter of formal notice, which marks the start of litigation.[20] There is a large literature on how the Commission chooses to open cases (Börzel 2001; Mbaye 2001; Jensen 2007; Sedelmeier 2008). I use a dataset of infringement cases from Fjelstul and Carrubba (2018), which covers 2003–2013. I update the dataset through 2015. It is important to note that just because the Commission brings a case against a member state, that does not mean that the case will reach the CJEU. Member states and the Commission can settle cases at any point in the infringement procedure (see Fjelstul and Carrubba 2018).

1.4.2 Measuring Preference Divergence

I develop a behavioral indicator of preference divergence between member states and EU policy-makers that is specific to harmonizing directives: how frequently member states use technical regulations as *de facto* trade barriers in non-harmonized sectors. By standardizing technical regulations across member states at the EU-level, harmonization prevents member states from using technical regulations as *de facto* trade barriers. Thus, the more frequently a member state attempts to use technical regulations to discriminate against firms from other EU member states in non-harmonized industries, the more resistant we should expect that member state to be to implementing new harmonizing directives that take away that tool, which implies a higher level of preference divergence.

[19] Harmonization directives are marked with the subject matter code "approximation of laws."

[20] There are three types of infringement cases: non-communication cases, non-conformity cases, and bad application cases. I only consider non-communication cases, which the Commission initiates when member states fail to pass domestic legislation that transposes directives by the specified deadline.

In non-harmonized industries, governments often use technical regulations as *de facto* trade barriers to protect domestic firms (in violation of the EU Treaties).[21] By imposing a new technical regulation that domestic firms already comply with, a government can discriminate against foreign firms without imposing tariffs or quantitative restrictions (which are easier for the Commission to identify). Since detecting discriminatory regulations requires substantial industry-specific knowledge, technical regulations are a popular way to protect domestic firms. Noncompliance in the implementation of harmonizing directives provides domestic firms in the newly harmonized industry with a competitive advantage. If domestic firms are spared from having to comply with new regulations (which can be very costly), while firms in other EU member states are not, those domestic firms have an opportunity to expand market share, gaining a competitive advantage.

I leverage public reporting requirements to identify attempts by member states to use technical regulations as *de facto* trade barriers.[22] Before a member state government can adopt a proposed technical regulation, it must notify the Commission under a monitoring procedure called the 2015/1535 procedure.[23] Governments must do this early enough in the policy-making process that the notified technical regulation can still be substantively amended. The Commission uses a public database called the Technical Regulation Information System (TRIS) to track notified technical regulations. Over the period I am considering (2004–2015), there are 8,090 notified technical regulations. There is an automatic 3 month standstill period in which the notifying member state cannot adopt the notified technical regulation.[24] If a third-party member state believes that a notified technical regulation would

[21] The EU Treaties prohibit member states from creating *de facto* trade barriers that discriminate against firms from other EU member states (Articles 34–36 TFEU).

[22] The CJEU ruled in *CIA-Security* that if a member state enacts a technical regulation without notifying the Commission under the 2015/1535 procedure, national courts can invalidate the regulation.

[23] The 2015/1535 procedure is named after the most recent directive that amended it, Directive (EU) 2015/1535. The procedure was originally created by Council Directive 83/189/EEC and was previously amended by Directives 98/34/EC and 98/48/EC.

[24] The Commission can block a regulation for 12 months if it is working on a directive that will harmonize the industry. The standstill can be extended to 18 months if the Council adopts a common position during the original 12 month standstill period.

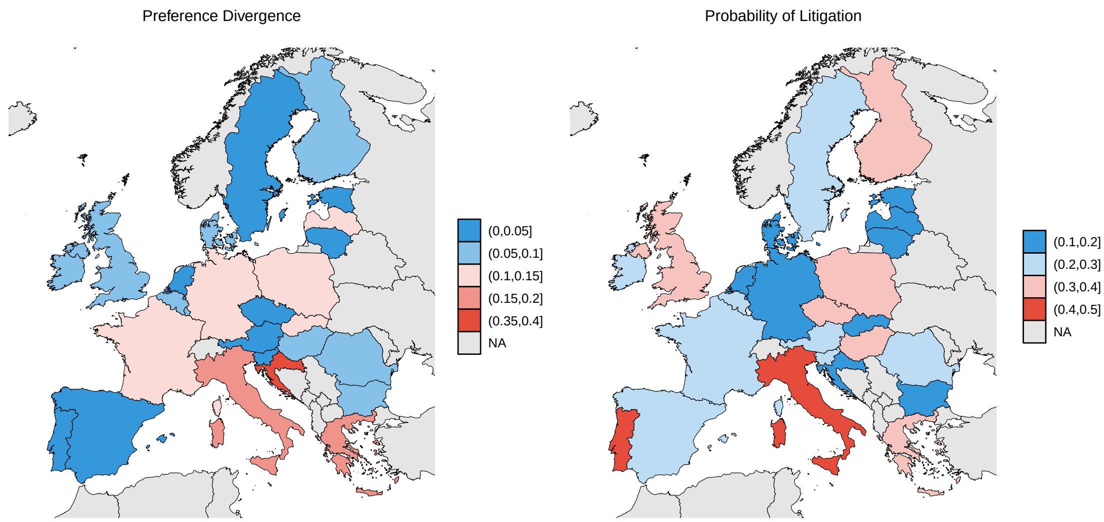

Figure 1.3. Geographical Distribution of Independent and Dependent Variables

Note: This figure shows the average level of preference divergence (the independent variable of interest) and the unconditional probability of an infringement case (the dependent variable) by member state across the entire sample (2004–2015).

constitute a *de facto* trade barrier, it can submit a *detailed opinion* (a formal complaint) to the Commission, which extends the standstill to 6 months.[25]

We can interpret a detailed opinion as an allegation by a third-party member state that the notifying member state is attempting to use a notified technical regulation as a *de facto* trade barrier.[26] Detailed opinions are credible allegations because they prohibit the parliament of the notifying member state from adopting the notified technical regulation for 6 months, which is a significant restriction on sovereignty.

Using the TRIS database, I calculate the average number of detailed opinions per notified technical regulation per member state (the notifying member state) per year. This is a time-varying measure of how much push-back a member state is getting from third-party member states in response to any new technical regulations it is trying to enact. Higher values indicate a preference for using technical regulations as *de facto* trade barriers, and

[25] The CJEU ruled in *Unilever* that if the notifying member state does not respect the standstill, national courts can invalidate the technical regulation.

[26] Opinions only indicate opposition to proposals, not support.

therefore higher preference divergence between the member state and EU policy-makers. If a member state has notified fewer than 3 technical regulations in a year, I code the variable as missing (there is too little information to code the variable). This reduces the sample from 14,747 observations to 13,413 observations.[27] Figure 1.3 shows the average level of preference divergence and the unconditional probability of an infringement case by member state across the entire sample (2004–2015).

In sum, the average number of detailed opinions per notified technical regulation captures a preference for using technical regulations as *de facto* trade barriers. Since harmonizing directives eliminate the opportunity to use technical regulations in this manner, this measure captures the divergence in the policy preferences of individual member state governments and the EU policy-making institutions.

1.4.3 Estimation and Analysis

To test Result 2, I estimate Bayesian multilevel logit models.[28] Observations are not independent because implementation opportunities are nested within member states. Thus, the errors for observations involving the same member state are likely to be correlated. Multilevel models explicitly account for this nested structure. I estimate varying-intercept multilevel logit models, where the intercept varies by member state.

In a Bayesian framework, the data are fixed and the parameters have a distribution. We start with prior beliefs about the parameters and use the data to update our beliefs,

[27] An industry-specific measure would be ideal, but it is not possible to code the industries that notified technical regulations apply to with available data. The TRIS database does not provide enough information about the specific content of notified technical regulations to cleanly map them to directives that harmonize technical regulations in the same industries.

[28] I estimate all models using `Stan` via `rstanarm` in `R` (Carpenter et al. 2017), which performs Markov chain Monte Carlo (MCMC) sampling of the posterior distribution using the NUTS algorithm (instead of the Metropolis-Hastings algorithm). Fewer MCMC samples are needed because the algorithm provides better coverage of the posterior distribution (Hoffman and Gelman 2014). I estimate 4 MCMC chains with 2,000 iterations each and a burn-in period of 1,000 iterations, producing a sample of 4,000 draws.

resulting in a posterior distribution. I use weakly informative priors.[29] Trace plots indicate that the models converge. The MCMC chains are stationary and there is high mixing. Autocorrelation plots show low autocorrelation for all parameters. Frequentist multilevel models yield similar results.

I control for three potentially confounding attributes of harmonizing directives. First, I control for whether the directive was enacted by the Council or the Commission.[30] Second, I control for the relative influence of the member state in the legislative process using its qualified majority voting (QMV) weight.[31] I normalize the QMV weights based on the number of member states so that it captures the relative influence of each member state over time. Third, I control for the number of cases in the Commission's docket. This captures the opportunity cost of bringing a case.[32]

Figure 1.4 shows the posterior distributions of the independent variables. It includes the mean of the posterior distribution for each covariate (the points), 50 percent Bayesian credible intervals (the thick lines), and 95 percent Bayesian credible intervals (the thin lines). There is an 95 percent chance that the parameter effect is in the 95 percent credible interval. In Models 1 and 2, I measure preference divergence as the mean number of detailed opinions per notified technical regulation. In Models 3 and 4, I use an alternative measure: the proportion of notified technical regulations where at least one member state files a detailed

---

[29] I prefer weakly informative priors to flat priors, which put too much probability on extreme values (Gelman 2006; Gelman and Hill 2006).

[30] Commission directives are enacted using delegated powers, so there is room for the Commission to enact policies that member states disagree with, which would result in higher preference divergence on average. Alternatively, member states might not have well-formed preferences over Commission directives, which are more technical and less politically salient than Council directives. This would result in lower preference divergence on average. The Commission could prioritize bringing infringement cases for Council directives, due to their higher political salience, but they could also prioritize bringing cases for Commission directives because there is a better chance they will be able to correct noncompliance.

[31] Member states with a large QMV weight are less likely to be outvoted in the Council, which means their preference divergence should be lower on average. Their influence in the Council could make it easier for them to stand up to the Commission, which may deter the Commission from bringing a case. At the same time, though, bringing these member states into compliance is a bigger win for the Commission.

[32] I standardize the QMV weight and docket size variables to have a mean of 0 and a standard deviation of 1 to help the models converge and to aid interpretation.

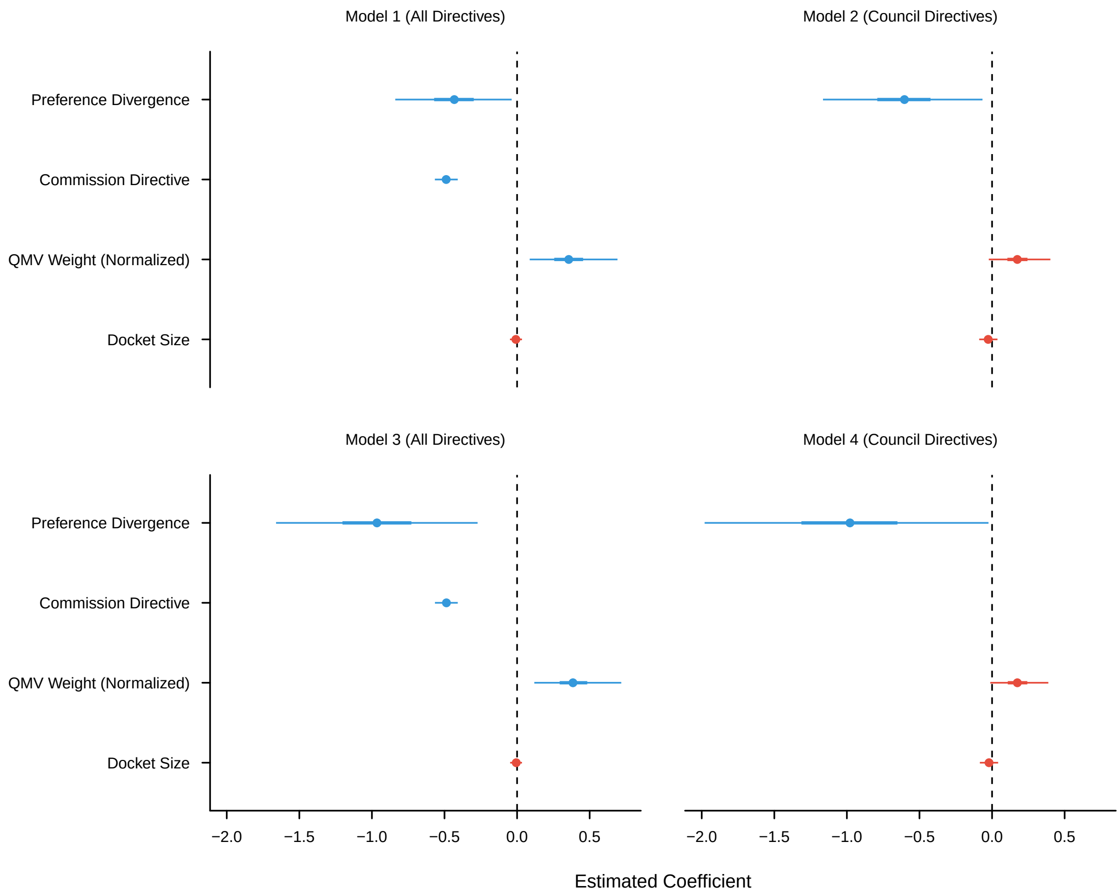

Figure 1.4. Coefficient Plot for Bayesian Multi-Level Models

Note: This figure shows the estimated posterior distribution for each independent variable. The points indicate the means. The thick bars indicate 50 percent Bayesian credible intervals and the thin bars indicate 95 percent Bayesian credible intervals. The 95 percent intervals for preference divergence do not overlap 0 in any of the models.

opinion. Models 1 and 3 use the full sample (13,413 observations), whereas Models 2 and 4 only include implementation opportunities involving Council directives, which are more politically salient than Commission directives (5,400 observations).

In all four models, preference divergence has a negative effect on the probability that the Commission opens an infringement case. This is consistent with my expectation that the data is likely to come from the region of the equilibrium space where the effect of preference divergence on the probability of litigation should be negative. This finding is consistent with the model, but inconsistent with the predictions of a model that does not take into account the strategic behavior of the court. The credible intervals for the measures

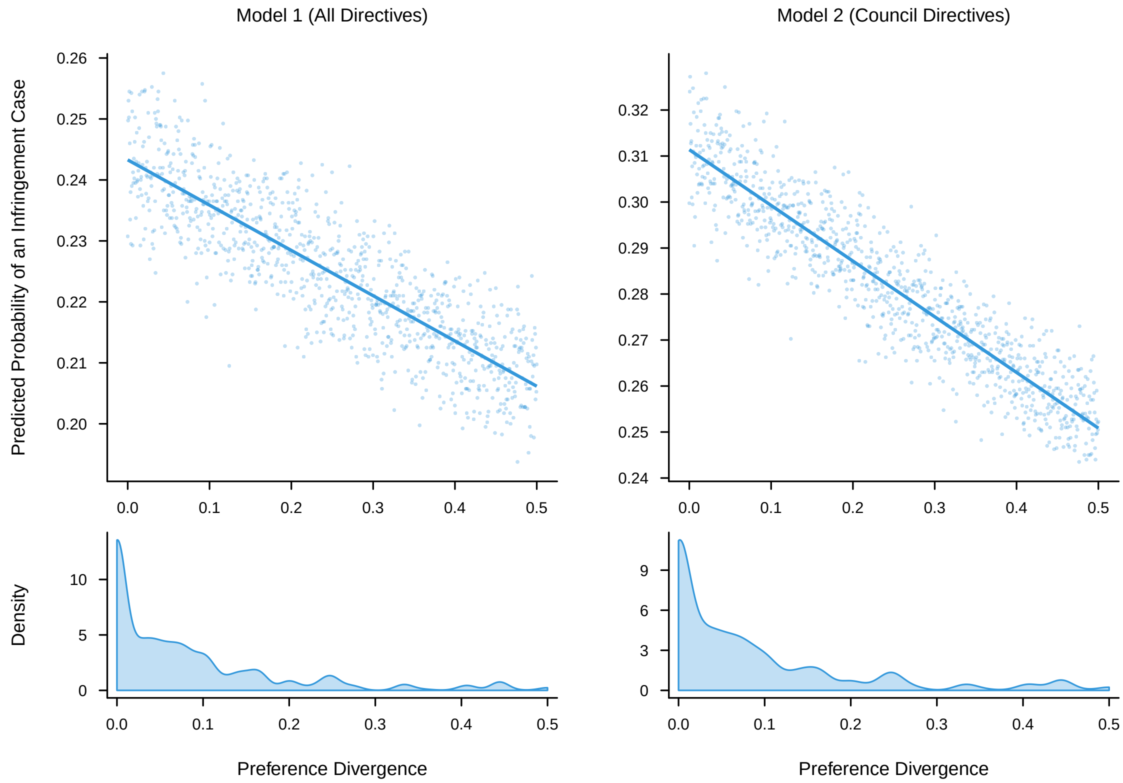

Figure 1.5. Marginal Effect of Preference Divergence

Note: This figure shows the effect of preference divergence on the probability that the Commission opens an infringement case based on Models 1 and 2. Model 1 includes all directives and Model 2 includes only Council directives. In Model 1, the predicted probability of a case decreases from approximately 24 percent to 21 percent. In model 2, it decreases from approximately 31 percent to 25 percent.

of preference divergence exclude 0, indicating that there is a greater than 95 percent chance that increasing preference divergence decreases the likelihood of an infringement case.

To understand the substantive significance of the results, I plot marginal effects. First, I simulate data, varying preference divergence (from the in-sample minimum to the in-sample maximum) while holding the other covariates at their means.[33] Second, I draw 4,000 outcomes (0 or 1) from the posterior predictive distribution (4,000 is the size of the posterior sample, which is the combined length of the four MCMC chains after discarding the burn-in period) conditional on the simulated data. Thus, there are 4,000 draws from the posterior predictive distribution for each simulated value of preference divergence. Third, I plot the

[33] I do not condition on the member state-specific intercepts.

predicted probability of the case as a function of preference divergence by taking the mean of the 4,000 simulated outcomes at each simulated value of preference divergence. I also include a best-fit line.

Figure 1.5 shows the average marginal effects for Models 1 and 2. In Model 1, the predicted probability of an infringement case decreases from approximately 24 percent to 21 percent. In Model 2, which only includes Council directives, it decreases from approximately 31 percent to 25 percent, a 20 percent decrease.[34] It makes sense that the average marginal effect is larger in Model 2, as Council directives are more politically salient than Commission directives; consequently, the CJEU is more likely to care about *ex post* compliance, which should induce a larger change in the behavior of the Commission. The estimated marginal effects for Models 3 and 4 are substantively similar.

Using Model 2 (Council directives only), I estimate the average marginal effect for each member state, conditioning on the member state-specific intercepts, setting the docket size variable at its global in-sample mean, and setting the QMV weight variable at its within-member state mean. Note that because Model 2 is a non-linear model, the effect of preference divergence depends on the values of the covariates, including the member state-specific intercepts. Figure 1.6 shows the predicted change in the probability that the Commission opens an infringement case associated with varying preference divergence from its global in-sample minimum to its global in-sample maximum. The effect varies from a decrease of approximately 4.75 percentage points to a decrease of approximately 7.25 percentage points. If we assume that the model captures the data generating process, we can infer that the Court is more sensitive to *ex post* compliance by member states for which the estimated effect is larger, such as the United Kingdom, Italy, and Greece.

If we are willing to assume that my formal model captures the data generating process, these findings imply that the CJEU is sensitive to *ex post* noncompliance with adverse rulings

[34] The effects are similar for my other measure of preference divergence, which is the proportion of notified technical regulations per year in which at least one detailed opinion is filed.

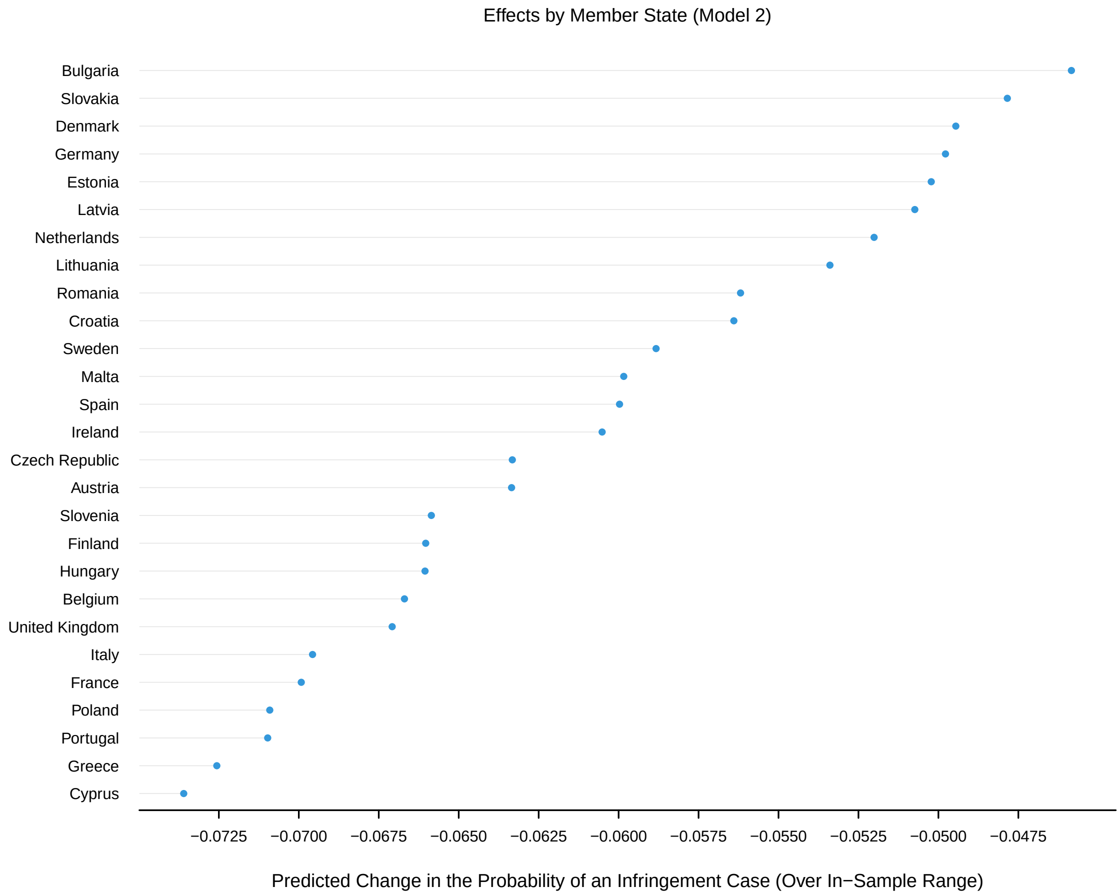

Figure 1.6. Marginal Effect of Preference Divergence by Member State

Note: This figure shows the predicted change in the probability of an infringement case associated with a change in preference divergence from its global in-sample minimum to its global in-sample maximum. (based on Model 2). The size of the effects varies from a decrease of approximately 4.75 percentage points to a decrease of 7.25 percentage points.

(Assumption 1). Consistent with Carrubba and Gabel (2015), the preferences of member states constrain the ability of the CJEU to correct noncompliance. These findings also imply a compliance deficit in the EU due to strategic noncompliance, which is consistent with recent empirical work (König and Mäder 2014; Fjelstul and Carrubba 2018).

1.5 Conclusion

This paper contributes to the literatures on courts and policy implementation by analyzing how the politics of adjudication affect the implementation of policy. Recent research shows that constitutional courts and quasi-constitutional international courts anticipate noncompliance with their rulings and are hesitant to rule against noncompliant defendants when it is costly to comply with an adverse ruling (e.g., Carrubba and Gabel 2015).

My model predicts that implementing actors will more closely implement laws for intermediate levels of preference divergence between the implementing actor and a policy-maker. For low values of preference divergence, the plaintiff is unlikely to bring a case because the benefits of correcting minimal noncompliance do not justify the costs of litigating. For high values of preference divergence, on the other hand, the plaintiff is unlikely to bring a case because the court, anticipating noncompliance with an adverse ruling, is unlikely to rule against the defendant. Looking at the Court of Justice of the European Union (CJEU), I find empirical evidence that is consistent with the model.

This study expands our understanding of how courts affect the behavior of implementing actors. I find that implementing actors try to thread a needle: they choose to implement a policy that is not fully compliant, but complaint enough that the likelihood of litigation is not too high. They trade a worse policy outcome for a lower chance of going to court. Consequently, none of the actors get their first-best policy outcome; implementing actors make concessions to avoid litigation, and plaintiffs are deterred from litigating the most severe instances of noncompliance, which means that policy-makers have to settle for partially correct implementation in expectation.

This study also has important implications for our understanding of the conditions under which international courts, like the CJEU, are effective at facilitating international cooperation. If courts are not sensitive to noncompliance with their rulings (*ex post* noncompliance), then when the cost of compliance is high, member states could prefer to exit rather than complying with adverse rulings or paying the political costs of ignoring them (Johns 2015). However, if they are sensitive to *ex post* noncompliance, then international courts facilitate the stability of international agreements at the expense of higher levels of noncompliance. When the cost of compliance is high, the court is not effective at deterring noncompliance, so member states can get away with implementing policies that are more consistent with their preferences; thus, they do not have an incentive to exit the regime.

Future research on courts and policy implementation should endogenize the policy-making process (i.e., the legislative process by which policy-makers bargain over *de jure* policies) and look for evidence of the discriminating predictions of the model in other empirical contexts. Future research should also consider how the internal judicial politics of constitutional and international courts affects their ability to strategically anticipate noncompliance with adverse rulings, as that strategic behavior drives their effectiveness at deterring noncompliance by implementing actors.

Paper 2:

The Political Economy of Noncompliance in Common Markets

Abstract

States create common markets to accrue consumer welfare gains. Given incentives to cheat to protect domestic firms from foreign competition, they create international regulatory regimes to manage noncompliance. I develop a formal model that explains how the politics of compliance in regulatory regimes systematically distorts the welfare gains that states accrue from developing common markets. The model predicts that regulatory regimes are most effective at enforcing compliance (i.e., at reducing trade barriers) in sectors with intermediate levels of firm homogeneity in terms of productivity. In highly homogenous sectors, regulatory regimes are not effective because noncompliance is minimal enough that litigation is not cost-effective; in highly heterogenous sectors, regulatory regimes are not effective because courts, concerned about noncompliance with their rulings, are unlikely to rule against defendants, deterring plaintiffs from bringing cases. The model also predicts the downstream consequences for the performance of individual firms and consumer welfare.

2.1 Introduction

States create common markets — a form of economic integration in which a regional trade bloc eliminates internal trade barriers — to improve consumer welfare.[1] Common markets improve consumer welfare by altering the composition of the economy: productive, exporting firms gain market share at the expense of unproductive, import-completing firms, causing prices to drop (Melitz 2003; Chaney 2008).[2] However, governments also have political incentives to cheat by imposing discriminatory trade barriers that protect those unproductive, import-completing firms from foreign competition, reducing the impact of the common market on their profitability. Given the incentives for noncompliance, under what conditions do common markets actually generate welfare gains?

To answer this question, we need to take into account the politics of compliance. Aware of the incentives to protect domestic firms, states rationally design international regulatory regimes to manage noncompliance. In particular, they create international courts to adjudicate disputes over noncompliance. Common markets generate welfare gains when regulatory regimes are effective at enforcing compliance with rules of the market. But these regimes only work when litigants are willing to bring cases and courts are willing to rule against defendants. In practice, there is systematic bias in the noncompliance cases that are litigated (König and Mäder 2014; Fjelstul and Carrubba 2018).

In this paper, I show that common markets are most welfare-enhancing in sectors of the economy where firms are highly heterogenous in terms of their productivity; this is when states are most tempted to erect trade barriers, and therefore when an agreement to eliminate trade barriers does the most to enhance consumer welfare. However, this is also when regulatory regimes are least effective at enforcing compliance (i.e., at reducing trade

[1] Trade barriers include tariffs and non-tariff barriers, such as quantitative restrictions and product standards that *de facto* discriminate against foreign goods.

[2] Productive firms have lower marginal costs, and they pass on these savings to consumers in the form of lower prices.

barriers), undermining the productivity gains that states actually accrue in practice. When the potential economic gains are greatest, the politics are most pernicious.

To identify the economic conditions under which common markets improve consumer welfare, given the politics of compliance in regulatory regimes, I embed a model of international trade (with firms and consumers) in a model of compliance (with governments, litigants, and a court). I do this by micro-founding the costs of compliance in the economy of the trade model. This allows the costs of compliance to be a function of the distributive consequences of trade liberalization: productive, exporting firms gain market share at the expense of unproductive, import-competing firms. I model governments as having politically motivated preferences over economic outcomes. The trade barriers they choose to impose can then affect those outcomes.

The model predicts the sectors of the economy in which regulatory regimes will be effective at reducing trade barriers as well as the downstream consequences for the performance of individual firms and for consumer welfare gains. Regulatory regimes are most effective at reducing trade barriers in sectors with intermediate levels of firm homogeneity. These are therefore the sectors in which the distributive consequences of a common market are highest and in which member states accrue the largest welfare gains. In highly homogenous sectors, regulatory regimes are not effective because noncompliance is minimal enough that litigation is not cost-effective; in highly heterogenous sectors, regulatory regimes are not effective because courts, concerned about noncompliance with their rulings, are unlikely to rule against defendants, deterring plaintiffs from bringing cases.

The question of how common markets actually affect the political economy of regional blocs is critical to our understanding of the modern global economy. Given the on-going deadlock in World Trade Organization (WTO) negotiations, states have increasingly turned to regional integration as an alternative, which has resulted in a proliferation of common markets. The success of the European Union (EU) — the world's second largest economy — has also encouraged other regional blocs to develop their own common markets. Currently, 79 states are members of a (semi-)functional common market. Another 42 states are members of regional economic organizations that have announced plans to develop a common market (see Figure 2.1).

In sum, this paper contributes to the literature on international institutions by developing a theoretical account of how the politics of compliance in international regulatory regimes affects the original economic objective — the development of a complete common market — that the regime was created to achieve. The model predicts the types of sectors — those with intermediate levels of homogeneity — in which regulatory regimes will be more effective. It also predicts the downstream consequences for the performance of individual firms and consumer welfare gains.

2.2 International Regulatory Regimes

Noncompliance with the rules of common markets (i.e., the imposition of trade barriers) is very common, even in the EU (König and Mäder 2014; Fjelstul and Carrubba 2018). To manage noncompliance, states rationally design international regulatory regimes to adjudicate disputes over compliance with the rules of the common market (Koremenos, Lipson and Snidal 2001; Carrubba and Gabel 2015). States create bureaucracies to monitor and prosecute member state noncompliance and international courts to adjudicate disputes over noncompliance. (In some regimes, private actors can also bring noncompliance cases.) Once states create courts, there are two aspects to compliance. There is initial compliance with the rules of the regime (*ex ante* compliance), and there is compliance with the rulings of courts in noncompliance cases (*ex post* compliance). There is no guarantee that member states will respect adverse rulings (Garrett, Kelemen and Schulz 1998; Alter 2000; Conant 2002; Slepcevic 2009; Panke 2010; Carrubba and Gabel 2015).

Member state governments, litigants, and international courts all operate strategically within the formal noncompliance procedures of regulatory regimes. Their incentives produce

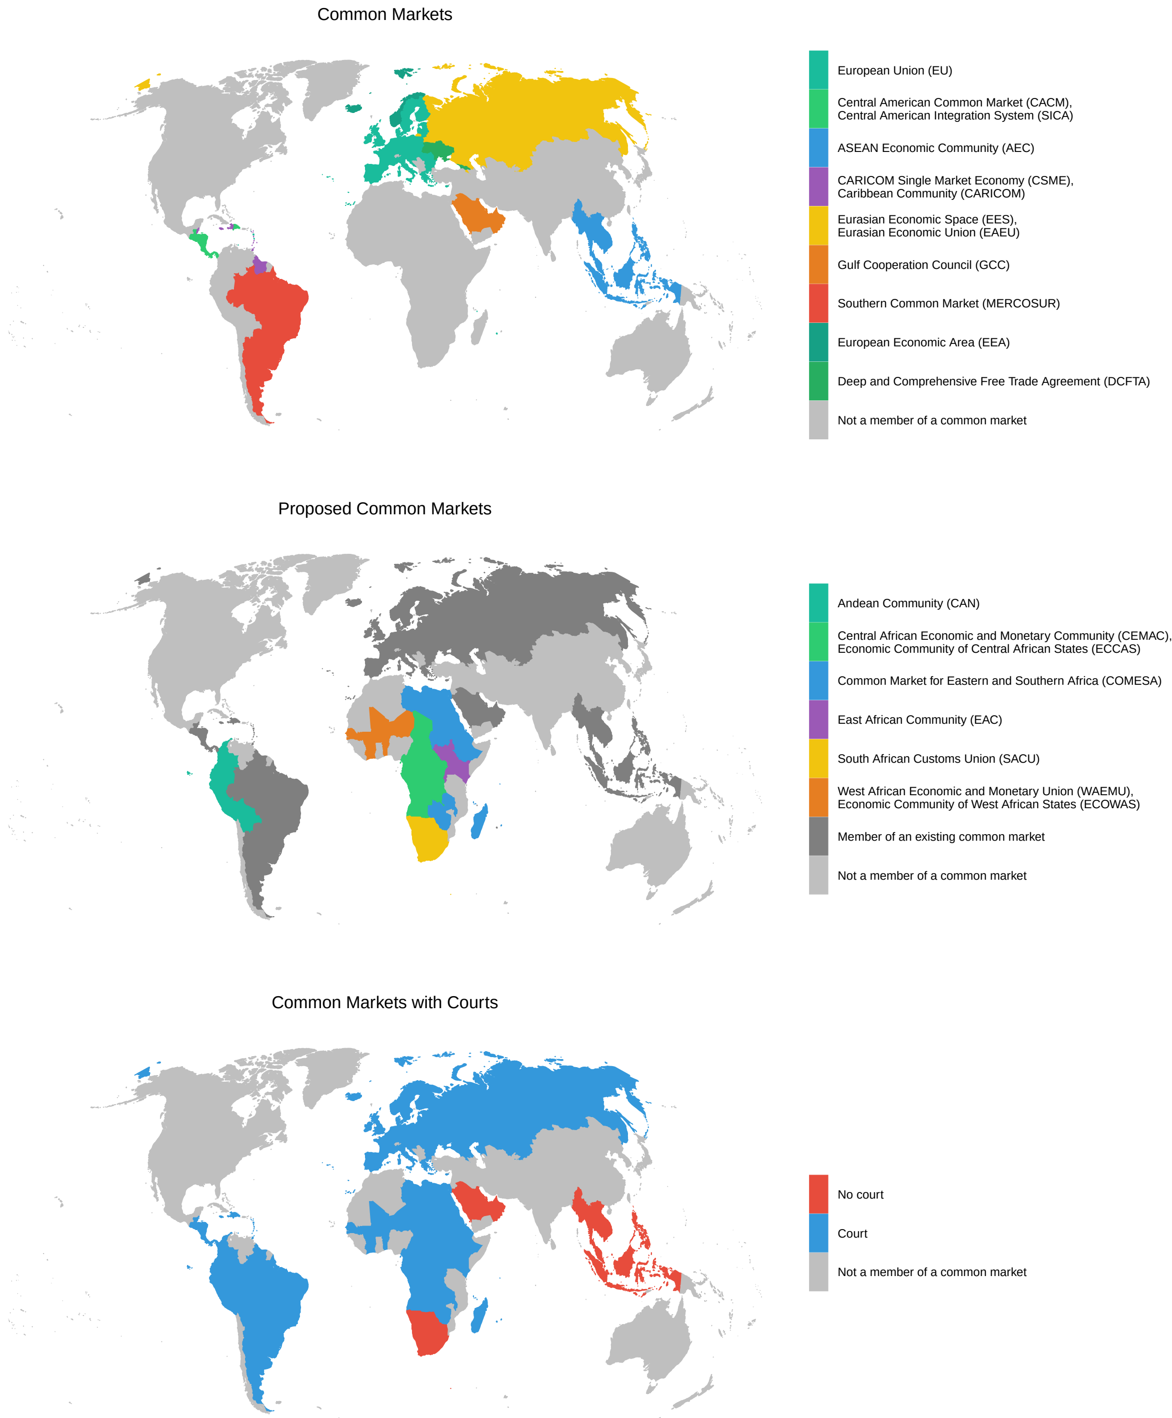

Figure 2.1. Common Markets

Note: The top map shows existing common markets (some are more complete than others). The middle map shows proposed common markets. The bottom map shows which of these existing and proposed common markets have an international court to adjudicate disputes over noncompliance.

a political process in which regulatory regimes successfully prevent or correct some violations, but permit others (Carrubba and Gabel 2015). Moreover, the politics of compliance generates systematic bias in the types of noncompliance cases that get litigated. Existing literature on compliance provides some general intuition about when noncompliance should get litigated: it depends on governments' costs of compliance (e.g., Carrubba and Gabel 2015; Fjelstul and Carrubba 2018).

Courts are concerned with *ex post* compliance with their rulings, and are therefore less likely to rule against governments when the costs of compliance are high (Alter 2000; Pollack 2003; Vanberg 2005; Carrubba 2005; Carrubba, Gabel and Hankla 2008; Carrubba 2009; Gilligan, Johns and Rosendorff 2010; Carrubba et al. 2012; Johns 2012; Carrubba and Gabel 2015; Martinsen 2015; Larsson and Naurin 2016). Litigants anticipate this behavior and drop cases when they are unlikely to win. The literature on international bureaucracies, like the European Commission, has long suspected that institutional litigants strategically choose which noncompliance cases to pursue (Mbaye 2001; Börzel 2003; Thomson, Torenvlied and Arregui 2007; Hartlapp and Falkner 2009; Steunenberg and Rhinard 2010). The most recent literature finds empirical evidence that the Commission drops cases when the costs of compliance are high (König and Mäder 2014; Fjelstul and Carrubba 2018). The key takeaway is that the politics of compliance leads to uneven enforcement outcomes, and that this variation is not random.

Since this bias in which cases are litigated is driven by the costs of compliance, we should expect the value of a regulatory regime to depend on the character of those costs. But without a theoretical model that explains where the costs of compliance come from, we cannot characterize how this bias will affect the ability of the regime to facilitate deep cooperation — the degree to which signing an international agreement causes states to behave differently than they would have otherwise (Downs, Rocke and Barsoom 1996).[3]

[3] Scholars have studied the conditions under which international institutions can facilitate deep cooperation across a wide variety of contexts (Keohane 1984; Chayes and Chayes 1993; Burley and Mattli 1993; Alter 2001; Rosendorff and Milner 2001; Stone Sweet and Brunell 1998; Rosendorff 2005; Simmons 2009).

2.3 Formal Model

I build up my formal model in three steps. I start with an open economy based on Melitz (2003), which is the starting point for most new-new trade theory (NNTT) models (Chaney 2008; Demidova and Rodríguez-Clare 2009; Melitz and Redding 2014; but not Melitz and Ottaviano 2008). Then, I add a policy-making subgame at the start of the game in which governments can choose trade barriers. This model serves as a counterfactual — it identifies the trade barriers that governments would choose in the absence of a regulatory regime. Next, I add a regulatory regime by adding a litigation subgame between the policy-making subgame and the economy subgame. After governments choose trade barriers, litigants can bring noncompliance cases, which a reduced-form court adjudicates. I use comparative statics to identify the sectors of the economy in which regulatory regimes will be most effective at reducing trade barriers relative to the counterfacutal. Then, I identify the downstream effects on firm performance and consumer welfare.

I micro-found the costs of compliance in an economy, which allows them to depend on the consequences of trade liberalization. Governments have preferences over distributive outcomes, and they can affect those outcomes by changing trade barriers (like Rosendorff 2005; but unlike Carrubba and Gabel 2015; Fjelstul and Carrubba 2018; Johns 2012). Existing models of trade in economics almost always treat trade barriers as exogenous (Melitz 2003; Chaney 2008; an exception is Demidova and Rodríguez-Clare 2009), but I allow governments to choose trade barriers in continuous space. Unlike optimal tariff models from economics, governments are not social planers or welfare-maximizers (e.g., Demidova and Rodríguez-Clare 2009). They have competing, politically motivated preferences: they care about the performance of import-competing domestic firms (due to lobbying) and consumer welfare (due to electoral incentives), both of which depend on *ex ante* trade barriers and whether they survive litigation.

I base the economy on a standard new-new trade theory (NNTT) model (i.e., Melitz 2003; Chaney 2008; Melitz and Redding 2014), which more accurately captures the process by which firms select into exporting — a firm's productivity determines whether exporting is profitable — than classical models (Heckscher-Ohlin and Ricardo-Viner) or new trade theory (NTT) models (Krugman 1980). Unlike these other trade theories, NNTT correctly predicts that only the most productive firms export. Firm selection into exporting affects prices, which affects consumer welfare. Since my objective for the model is to identify how the politics of compliance distorts consumer welfare gains, it is therefore very important to model firm selection correctly.

2.4 An Open Regional Economy

I start by modeling a one-sector open regional economy with n symmetric countries. There are two types of actors: firms and consumers. There is a mass of firms M in each country. Each firm has a productivity $\varphi > 1$, which is drawn from a probability density function, $g(\varphi)$.[4] Each firm produces a unique variety of good $\omega \in \Omega$. Thus, while ω uniquely identifies a firm, multiple firms can have the same productivity φ. In equilibrium, all firms with the same productivity φ will behave identically. Each country has one representative consumer, with income I. Unlike Melitz (2003), I assume that the mass of firms M in each country and consumer income I are exogenous, which means that the economy is in partial equilibrium. Solving for a general equilibrium adds considerable complexity to the model without changing any of the results.[5]

The order of play is as follows. Firms choose whether to produce for the domestic market and whether to export to foreign markets. Conditional on serving a market, each firm

[4] The corresponding cumulative distribution function is $G(\varphi)$.

[5] A general equilibrium features free entry (firms choose whether to enter the market prior to learning their productivity and only enter when the expected profits exceed a fixed cost of entry) and labor market clearing (total firm revenue equals total labor payments). See Melitz (2003) for details.

chooses a price to charge, $p(\varphi)$. Then, the representative consumer in each country chooses a quantity of each available variety $q(\omega)$ to purchase.

2.4.1 Demand

The representative consumers have constant elasticity of substitution (CES) preferences (Dixit and Stiglitz 1977). This is the standard approach to modeling consumers in NTT and NNTT models (Krugman 1980; Melitz 2003; Chaney 2008; but not Melitz and Ottaviano 2008). The consumer demands at least some of each available variety.[6] In this sense, she has a love for variety. Her income I limits the quantity of available varieties she can buy. Each representative consumer's utility is:

$$u_c = \left[\int_{\Omega} q(\omega)^{\frac{\sigma-1}{\sigma}} \, \mathrm{d}\omega \right]^{\frac{\sigma}{\sigma-1}}, \tag{2.1}$$

where $q(\omega)$ is the quantity of each variety demanded (the choice variable), σ is the elasticity of substitution, and Ω is the set of varieties that are available.[7]

CES preferences introduce monopolistic competition. Under monopolistic competition, firms perceive competition from other firms, but pricing is not a strategic game between firms (unlike oligopolistic competition). Firms also have market power, which means that a firm can change consumer demand for its variety by changing the price that it charges. In equilibrium, firms can make a profit in the short run (unlike perfect competition, where firms do not make a profit). Thus, using CES preferences will allow me to study how international regulatory regimes will affect firm performance.

[6] Whether any given variety ω is available for purchase in a given market depends on whether the firm that produces it chooses to produce for that market.

[7] A high σ implies a weaker love of variety because small changes in price will cause a consumer to shift more of her consumption to cheaper varieties. As σ goes to infinity, varieties become perfect substitutes. As it goes to 0, varieties become perfect complements.

2.4.2 Supply

Firms pay a per-unit cost to produce a good. I assume there is only one factor of production and normalize the cost of that factor to one per unit.[8] More productive firms enjoy lower marginal costs. A firm's marginal cost in the domestic market is the inverse of its productivity, $c_d(\varphi) = \frac{1}{\varphi}$.

To produce any amount of goods for any market, a firm must pay a fixed cost f.[9] We can think of these fixed costs as marketing costs. Firms pay extra variable costs b to export due to trade barriers. In the next section, I will endogenize these trade barriers b by allowing governments to choose them. The marginal cost of selling one unit in a foreign market is the marginal cost multiplied by trade barriers, $c_x(\varphi) = \frac{b}{\varphi}$.

The net profit that a firm makes in a market is gross profit less fixed costs. Gross profit is the quantity sold $q(\omega)$ times per-unit profit, which is the price of a unit $p(\omega)$ minus the marginal cost $c(\omega)$ of producing it:

$$\pi(\omega) = q(\omega)\big(p(\omega) - c(\omega)\big) - f. \tag{2.2}$$

I assume that firm productivity is Pareto distributed (e.g., Chaney 2008). The empirical literature in economics finds that firm productivity is approximately Pareto distributed (Axtell 2001; Luttmer 2007; Helpman, Melitz and Yeaple 2004; Gabaix 2009). The probability density function (PDF) and the cumulative distribution function (CDF) for the Pareto

[8] In standard NNTT models, the only factor of production is labor. In Chaney (2008) and related models, wages are exogenous and pinned down by the presence of an outside sector.

[9] The fact that firms pay fixed costs to produce and export introduces increasing returns to scale (Krugman 1980), an innovation of new trade theory (NTT) models and a feature of all NNTT models. Increasing returns to scale account for why we observe intra-industry trade (i.e., trade flows between two countries within the same industry), which is not predicted by classical theories (e.g., the Heckscher-Ohlin and Ricardo-Viner models). Intra-industry trade is prevalent in common markets, which is another reason to base the economy in the model on NNTT instead of classical theories.

distribution are, respectively:

$$g(\varphi) \equiv \frac{\theta}{\varphi^{\theta}} \qquad \text{and} \qquad G(\varphi) \equiv 1 - \varphi^{-\theta}, \tag{2.3}$$

where θ is the shape parameter of the distribution.[10] A high θ means that firms are more homogeneous and a low θ means that firms are more heterogeneous. We can think of θ as capturing the structure of the sector in terms of firm productivity.

2.4.3 Open Economy Equilibrium

Proposition 4 summarizes equilibrium behavior in the economy.[11] Firms only produce for the domestic market if they are sufficiently productive: $\varphi > \varphi_d^*$. Similarly, firms only export if they are sufficiently productive: $\varphi > \varphi_x^*$. Due to the variable costs of trade, firms will only export if they sell to the domestic market: $\varphi_x^* > \varphi_d^*$. Conditional on selling to a market, firms choose an optimal price $p^*(\varphi)$, which is a constant markup over marginal cost: $p^*(\varphi) = \left(\frac{\sigma}{\sigma-1}\right) c(\varphi)$.

Proposition 4. The equilibrium of a one-sector open economy with symmetric countries and heterogeneous firms that produce substitutable varieties under monopolistic competition is:

1. Firms sell to the domestic market if they are sufficiently productive: $\varphi > \varphi_d^*$.
2. Firms export if they are sufficiently productive: $\varphi > \varphi_x^*$. Firms will only export if they sell to the domestic market: $\varphi_x^* > \varphi_d^*$.
3. Conditional on selling to a market, firms choose an optimal price $p^*(\varphi)$, which is a constant markup over marginal cost.
4. The representative consumer in each country chooses an optimal quantity $q^*(\varphi)$ of available varieties to consume, subject to an income constraint.

---

[10] We must assume that $\theta > \sigma - 1$ for average firm productivity in equilibrium to be finite. Melitz (2003) and Chaney (2008) make the same assumption.

[11] Throughout, I use an asterisk to indicate equilibrium quantities.

The consumer in each country observes the price of each available variety and chooses an optimal quantity of available varieties to consume, subject to an income constraint. The consumer buys $q^*(\varphi) = p^*(\varphi)^{-\sigma} I (P^*)^{\sigma-1}$ of each available variety, where P^* is the equilibrium Dixit-Stiglitz price index (Dixit and Stiglitz 1977; Chaney 2008). The price index is the cost of obtaining one unit of utility. See Appendix B for equilibrium equations and full proofs.[12]

2.4.4 Comparative Statics

To lay the groundwork for the remainder of the paper, I show how changes in trade barriers and sector homogeneity affect firm behavior — whether a firm produces or exports — and consumer welfare. To shift the productivity distribution such that firms become more homogeneous, I increase the shape parameter of the Pareto distribution, θ. This concentrates more of the density on low productivity firms and decreasing θ distributes the density more evenly across firms (see Figure 2.2).[13] Due to the complexity of the closed-form solutions for the productivity cut-points, I calculate comparative statics numerically using Monto Carlo simulations.[14] See Appendix B for details.

As trade barriers b decreases, the domestic production cut-point increases φ_d^* and the exporting cut-point φ_x^* decreases (see Figure 2.3, Panels A and B). This changes the composition of the sector: more firms can afford to export, and fewer firms can afford to produce for the domestic market. More productive firms charge lower prices (they have lower marginal costs and pass some of these savings on to consumers in the form of lower prices), so this decreases the price index P^*. This increases consumer welfare W^*, which is defined as income divided by the price index.

[12] Adding free entry, labor market clearing, and solving for a general equilibrium, makes the model significantly more complicated. This additional complication does not change the results.

[13] Note that changing θ to make firms more homogeneous necessarily also decreases average productivity. This is a consequence of using the Pareto distribution.

[14] It is well-known in the economics literature that deriving smooth comparative statics in models of international trade with monopolistic competition is difficult.

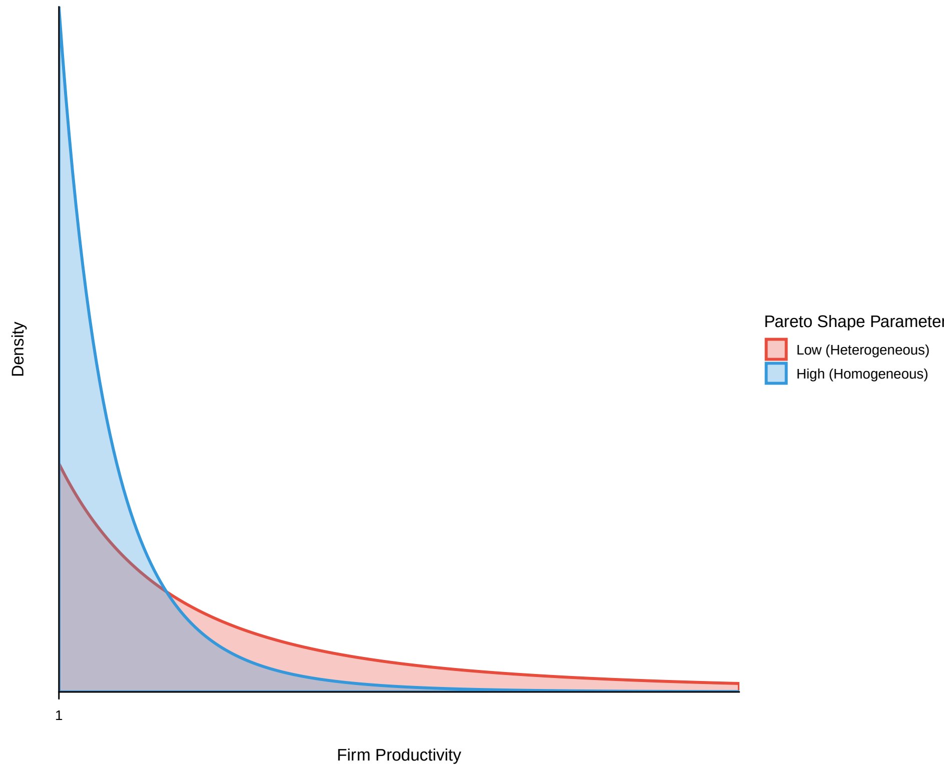

Figure 2.2. Distribution of Firm Productivity

Note: As the shape parameter of the Pareto distribution increases, firms become more homogeneous in terms of their productivity (and average productivity decreases).

As the sector becomes more homogeneous (as θ increases), both productivity cut-points decrease (see Figure 2.3, Panels C and D). In homogeneous sectors, price competition is less intense and market share is more evenly distributed across firms. This helps less productive firms cover the fixed costs of domestic production and exporting.

Result 6. As trade barriers b increase, the domestic production cut-point φ_d^* decreases and the exporting cut-point φ_x^* increases. As sector homogeneity θ increases, the productivity cut-points, φ_d^* and φ_x^*, decrease.

2.5 Adding Strategic Policy-Making

Next, I relax the assumption that trade barriers are exogenous by adding a policy-making subgame at the start of the game. In this subgame, symmetric, strategic governments with

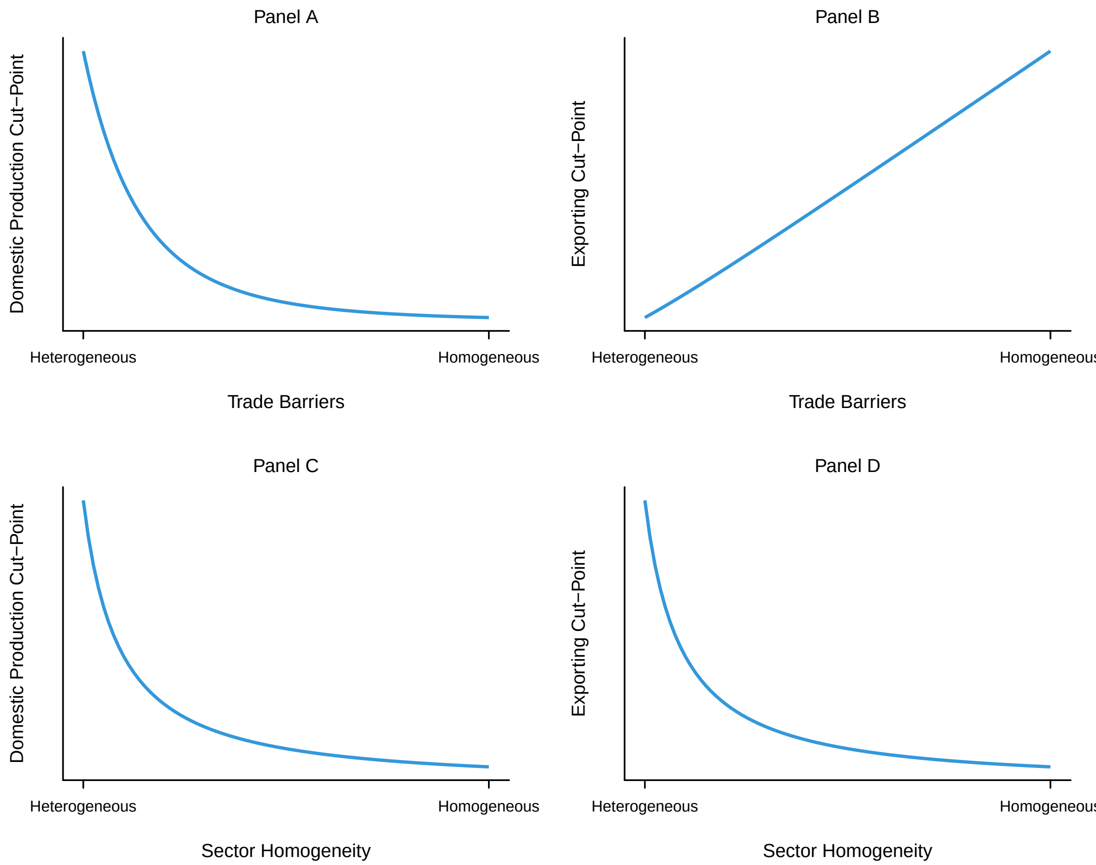

Figure 2.3. Comparative Statics in an Open Economy

Note: In an open economy, the domestic production cut-point is decreasing in trade barriers and the exporting cut-point is increasing in trade barriers. The domestic production and exporting cut-points are decreasing in sector homogeneity.

political preferences can choose optimal trade barriers $\tilde{b}^*$.[15] The economy then plays out just as before, conditional on $\tilde{b}^*$. This model serves as a counterfactual by establishing the trade barriers that governments would choose without a regulatory regime.

Governments balance competing interests in setting trade barriers. First, they want to protect at least some domestic firms from foreign competition. Trade barriers help domestic firms retain market share (Melitz 2003). This is especially important for non-exporters, which are less productive than exporters. As such, non-exporters stand to lose more market share to foreign competition than exporters. Domestic import-competing firms can often

[15] I use a tilde to denote equilibrium quantities in the counterfactual, without a regulatory regime.

overcome their collective action problem and lobby their government to protect them from foreign competition (Grossman and Helpman 1994).

Second, governments want to let in at least some foreign imports. Exporters are more productive than non-exporters, so they charge lower prices. Letting in imports also increases the number of unique varieties that are available in the domestic market. Importers pass on the extra costs of trade barriers to consumers in the form of higher prices (recall that prices are a constant markup over variable cost). Thus, from a consumer perspective, trade barriers are a tax. In addition, trade barriers increase the average prices charged by domestic firms by insulating unproductive firms (which charge higher prices) from competition. It is more difficult for consumers to overcome their collective action problem, but governments have a basic electoral incentive not to become too autarkic.

Governments want to choose trade barriers to optimally balance these competing producer and consumer interests. They want to strike a balance between lowering the domestic production cut-point to benefit domestic firms and lowering the exporting cut-point to benefit consumers. They cannot do both at the same time (see Result 6), as higher trade barriers b decrease the domestic production cut-point φ_d and increase the exporting cut-point φ_x. To introduce this tradeoff in the model, I make each government's utility decreasing in both the equilibrium domestic production cut-point φ_d^* and the equilibrium exporting cut-point φ_x^*, both of which are functions of trade barriers $\tilde{b}$:

$$u_g(\tilde{b}) = -w\varphi_d^* - (1-w)\varphi_x^*, \tag{2.4}$$

where $w \in (0,1)$ is the relative weight that governments place on the interests of domestic firms relative to the interests of domestic consumers (a large w implies a strong preference for protectionism). The government's utility function is decreasing linearly in both cut-points (but neither cut-point changes linearly with respect to trade barriers, as shown in Figure 2.3). Equation (2.4) produces well-behaved, single-peaked preferences.

In equilibrium, the government chooses optimal trade barriers $\tilde{b}^*$. There is a unique solution.[16] For intermediate values of w, there is an interior solution, $1 < \tilde{b}^* < \infty$. The optimal $\tilde{b}^*$ is the value at which the absolute marginal change in the domestic production cut-point $\tilde{\varphi}_d$ equals the absolute marginal change in the exporting cut-point $\tilde{\varphi}_d$. After the governments choose trade barriers, the economy plays out the same as before, conditional on optimal trade barriers $\tilde{b}^*$.

Proposition 5. Governments choose optimal trade barriers $\tilde{b}^*$ in equilibrium, where

$$\tilde{b}^* = \left[\frac{nw}{1-w}\right]^{\frac{1}{1+\theta}}.$$

Conditional on $\tilde{b}^*$, Proposition 4 describes firm and consumer behavior.

In equilibrium, as firms become more homogeneous, government prefer smaller trade barriers (see Figure 2.4).[17] In homogeneous sectors, price competition is less intense. Market share is more even distributed across firms, and more domestic firms can afford to stay in business. As such, there is less need for governments to protect domestic firms, which means they can reorient their trade policy towards consumer interests. By decreasing trade barriers, they encourage more foreign firms to enter the domestic market (the exporting cut-point is decreasing and the domestic production cut-point is increasing). However, the direct effect of changing sector homogeneity (see Result 6) overwhelms the effect of lower trade barriers. The net result is that both productivity cut-points are decreasing, allowing less productive firms to produce. Since unproductive firms charge higher prices, this decreases consumer welfare. This is the baseline against which I compare equilibrium behavior under a regulatory regime.

Result 7. The optimal trade barriers $\tilde{b}^*$ are decreasing in sector homogeneity θ. The productivity cut-points, $\tilde{\varphi}_d^*$ and $\tilde{\varphi}_x^*$, and consumer welfare $\tilde{W}^*$ are also all decreasing in sector homogeneity θ.

---

[16] See Appendix B for an analytical proof.

[17] I calculate comparative statics numerically using a Monte Carlo simulation. See Appendix B for details.

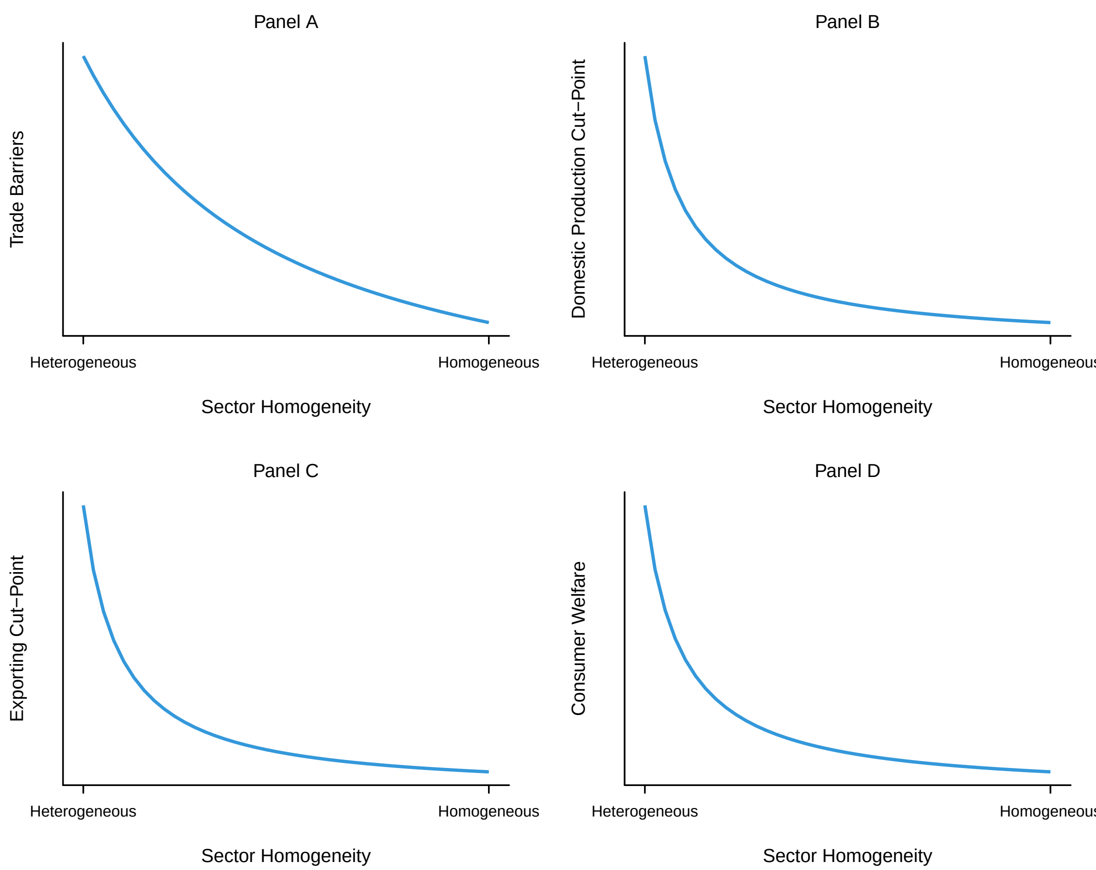

Figure 2.4. Comparative Statics with Strategic Policy-Making

Note: With strategic governments, optimal trade barriers, the domestic production cut-point, the exporting cut-point, and consumer welfare are all deceasing in sector homogeneity.

2.6 Adding an International Regulatory Regime

Next, I add an international regulatory regime to the model. A regulatory regime has three basic elements: (a) treaty that establishes a common market by requiring that all member states reduce their trade barriers to a particular level; (b) an international court that can adjudicate noncompliance cases against governments; and (c) a litigant that can initiate noncompliance cases. To model a regulatory regime, I add a litigation subgame between the policy-making subgame and the economy subgame. Now, after governments choose trade barriers, symmetric litigants in each country choose whether or not to bring a case, and a reduced-form court issues a ruling.

I model the treaty-mandated trade barriers as an exogenous parameter, $b_t \geq 1$. This parameter indicates the specific value of b that is fully compliant. It is common knowledge.[18] A value of $b_t = 1$ indicates that the common market does not permit any intra-regime trade barriers. Higher values indicate a higher tolerance for intra-regime trade barriers — and a less complete common market. Since the objective of common markets is to lower trade barriers, I consider a scenario in which the treaty-mandated trade barriers are lower (i.e., more liberal) than the optimal trade barriers the governments would choose in the counterfactual model: $b_t < \tilde{b}^*$.

Recent studies on international courts views courts as strategic actors (Vanberg 2015; Carrubba 2005; Carrubba and Gabel 2015, 2017). Courts care about the degree to which member states comply with their treaty obligations (*ex ante* compliance), but they also care about compliance with their rulings (*ex post* compliance). As such, a court anticipates how likely a government is to comply when it makes a ruling (Carrubba and Gabel 2015; Martinsen 2015; Larsson and Naurin 2016). A government is more likely to ignore the ruling of the court when the costs of compliance is high.

I incorporate this view of international courts by modeling the court as a reduced-form player: Conditional on the litigant bringing a case, there is some probability $h(c^*)$ of *ex post* compliance, which is the joint probability that the court rules against the government and that the government complies with the court's ruling. From the litigant's perspective, this is the probability of successful litigation. As Carrubba and Gabel (2015) show, this probability is endogenous to the government's cost of compliance in equilibrium c^*. As the cost of compliance increases, the conditional probability of *ex post* compliance decreases, $h'(c^*) < 0$.[19] The government is less likely to comply with an adverse ruling and, anticipating that, the court is more hesitant to rule against the government.

[18] Since government trade barriers b are also common knowledge, there is no uncertainty about whether governments have committed violations. This is similar to Fjelstul and Carrubba (2018), who do not allow a litigant to incorrectly infer noncompliance.

[19] I use a logistic function as the function form for $h(c^*)$. This converts c^* to a probability.

If the court rules in favor of the government, or if it rules against the government and the government does not comply with the adverse ruling, then the government can continue to keep the equilibrium *ex ante* trade barriers, $b_1^* = b_0^*$, that it has chosen. However, if the government loses and does comply *ex post* with the court's ruling, it comes into compliance with the treaty $b_1 = b_t$. Thus, equilibrium *ex post* trade barriers b_1^* are either the government's equilibrium *ex ante* trade barriers or the treaty-mandated trade barriers, $b_1^* \in \{b_0^*, b_t\}$, depending on how the litigation subgame plays out. In equilibrium, *ex ante* noncompliance is $\left| \, b_0^* - b_t \, \right|$ and *ex post* noncompliance is $\left| \, b_1^* - b_t \, \right|$.

The government's cost of compliance in equilibrium c^* is the absolute difference between the government's utility for fully complying the treaty, $u_g(b_t)$, and its equilibrium utility in the counterfactual, $u_g(\tilde{b}^*)$, where the government can choose an optimal trade barriers free from institutional constraints: $c^* = \left| \, u_g(b_t) - u_g(\tilde{b}^*) \, \right|$.[20] Note that the cost of compliance is always non-negative and that it is an equilibrium quantity because it depends on equilibrium behavior in the counterfactual model.

The new order of play is as follows. Symmetric governments in each country choose *ex ante* trade barriers b_0. Symmetric litigants in each country observe these trade barriers and choose whether or not to bring a case. If there is a case, there is some probability of *ex post* noncompliance $h(c^*)$, which determines *ex post* trade barriers. If there is *ex post* compliance, then the government comes into compliance with the treaty, $b_1 = b_t$; otherwise, it keeps the *ex ante* trade barriers it has already chosen, $b_1 = b_0$. If the litigant does not bring a case, *ex post* trade barriers also equal *ex ante* trade barriers, $b_1 = b_0$. From this point on, the economy subgame plays out exactly as before, conditional on the *ex post* trade barriers, $b_1 \in \{b_0, b_t\}$.

[20] The court does not want to issue a ruling that asks a government to come into compliance only for the government to ignore it. What matters, then, from the court's perspective, is how bad compliance is compared to what the government wants to do (not how bad compliance is compared to whatever policy that government chooses *ex ante*, which in equilibrium will be a compromise).

2.6.1 Litigant Preferences

The litigants care about *ex post* compliance with the treaty, $\left| b_1 - b_t \right|$. I assume that litigants prefer compliance: they suffer policy loss based on the difference between the treaty-mandated trade barriers b_t and *ex post* trade barriers, $b_1 \in \{b_0, b_t\}$. I use a simple exponential loss function to model this preference. Litigation is costly. If a litigant brings a case, it pays a cost k, where k is drawn from a distribution with commutative distribution function $J(k)$.[21] This cost is private information. The utility of the litigants is given by the following piece-wise function:

$$u_l(b_1) = \begin{cases} -(b_t - b_0)^2 & \text{if there is no case} \\ -(b_t - b_0)^2 - k & \text{if there is a case and no } \textit{ex post} \text{ compliance} \\ -k & \text{if there is a case and } \textit{ex post} \text{ compliance.} \end{cases} \tag{2.5}$$

If a litigant brings a case and wins, the government comes into compliance and implements the treaty-mandated trade barriers b_t, so the litigant does not suffer any policy loss and only pays the cost of bringing a case $-k$. If the litigant does not bring a case, or does bring a case, but there is *ex post* noncompliance (i.e., if the court rules in favor of the government or the government ignores an adverse ruling), its policy loss increases with the distance between the government's choice of trade barriers b_1 and the treaty-mandated trade barriers b_t.

2.6.2 Regulatory Regime Equilibrium

In equilibrium, litigants bring noncompliance cases when the cost of bringing a case is sufficiently small relative to the probability of *ex post* compliance, $k < k^*$ (similar to Fjelstul and Carrubba 2018). Otherwise, it does not bring a case. Thus, litigants drop cases when they are unlikely to win. This is consistent with recent empirical work on the European Commis-

[21] I use an exponential CDF as the functional form of $J(k)$. With this distribution, low-cost litigation is more likely than high-cost litigation.

sion, which finds evidence that the Commission drops costly cases (König and Mäder 2014; Fjelstul and Carrubba 2018), something earlier work had suspected (Mbaye 2001; Börzel 2003; Thomson, Torenvlied and Arregui 2007; Hartlapp and Falkner 2009; Steunenberg and Rhinard 2010). The probability that a litigant brings a case in equilibrium is the probability that a cost draw is below the cut-point: $\Pr(k < k^*)$ or $J(k^*)$.

Proposition 6. Under a regulatory regime, litigants bring noncompliance cases when the cost of bringing a case is sufficiently small: $k < k^* \equiv h(c^*)(b_t - b_0^*)^2$.

In equilibrium, governments anticipate the probability that the litigant in their country will bring a case, $J(k^*)$, and the conditional probability of an adverse court ruling, $h(c^*)$, and choose *ex ante* trade barriers b_0^* that maximize their expected utility, which is a function of *ex post* trade barriers, $b_1 \in \{b_0, b_t\}$:

$$E\big[u_g(b_0)\big] = J(k^*)\Big(h(c^*)u_g(b_t) + \big(1 - h(c^*)\big)u_g(b_0)\Big) + \big(1 - J(k^*)\big)u_g(b_0), \tag{2.6}$$

where $u_g(b_0)$ is given by Equation (2.4). Equation A.1 is a well-behaved, single-peaked concave function. There is a unique solution in which governments choose *ex ante* trade barriers that are optimal in expectation.[22] The economy plays out as before (see Proposition 4), conditional on *ex post* trade barriers, $b_1^* \in \{b_0^*, b_t\}$, which are stochastic. As such, we have to consider the economy in expectation, conditional on expected *ex post* trade barriers, $E\big[b_1^*\big]$. Expected *ex post* trade barriers are:

$$E\big[b_1^*\big] = J(k^*)\Big(h(c^*)b_t + \big(1 - h(c^*)\big)b_0^*\Big) + \big(1 - J(c^*)\big)b_0^*. \tag{2.7}$$

Proposition 7. Under a regulatory regime, governments anticipate the probability that litigants will bring a case, $J(k^*)$, and the probability of *ex post* compliance, $h(c^*)$, and

[22] I demonstrate this using a computational solution. See Appendix B for details.

choose optimal *ex ante* trade barriers, b_0^*. Conditional on expected *ex post* trade barriers, $E\left[b_1^*\right]$, Proposition 4 describes firm and consumer behavior.

2.6.3 Systematic Bias in Noncompliance Cases

I use this equilibrium to show how the politics of compliance in regulatory regimes generate systematic bias in the types of noncompliance cases that get litigated, and that this bias creates a distortion in the economy: regulatory regimes reduce trade barriers most in sectors with intermediate levels of homogeneity. Then, in the next section, I identify the downstream consequences for firm performance and consumer welfare.

I start by calculating the effect of a regulatory regime on *ex ante* trade barriers and *ex post* trade barriers in expectation, the latter being the ones that ultimately affect firms and consumers. To review, equilibrium *ex ante* trade barriers are b_0^*, expected equilibrium *ex post* trade barriers are $E\left[b_1^*\right]$, and equilibrium trade barriers in the counterfactual are $\tilde{b}^*$. The effect of the regime on *ex ante* trade barriers is the difference between equilibrium *ex ante* trade barriers and equilibrium trade barriers in the counterfactual: $b_0^* - \tilde{b}^*$. Similarly, the effect of the regime on expected *ex post* trade barriers is $E\left[b_1^*\right] - \tilde{b}^*$. *Ex ante* noncompliance is $\left|\, b_0^* - b_t \,\right|$ and expected *ex post* noncompliance is $\left|\, E\left[b_1^*\right] - b_t \,\right|$.

In equilibrium, a regulatory regime causes member state governments to reduce *ex ante* trade barriers in equilibrium, $b_0^* - \tilde{b}^* < 0$.[23] Governments make a concession to the regime in order to lower the probability of enforced compliance, which is the joint probability of a case and *ex post* compliance, $J(k^*)h(c^*)$. Figure 2.5 gives an example of an optimal concession. In equilibrium, *ex ante* trade barriers in expectation are less than optimal trade barriers in the counterfactual, but higher than the treaty-mandated trade barriers: $b_t < b_0^* < \tilde{b}^*$.

[23] I use Monte Carlo simulations for this and all subsequent results. I randomly draw values for exogenous parameters using probability distributions, solve the model numerically, and then calculate the effect of the regime on equilibrium quantities (relative to the counterfactual). I develop a computational algorithm to sign monotonic and non-monotonic comparative statics for sets of parameter values that yield interior solutions. See Appendix B for details.

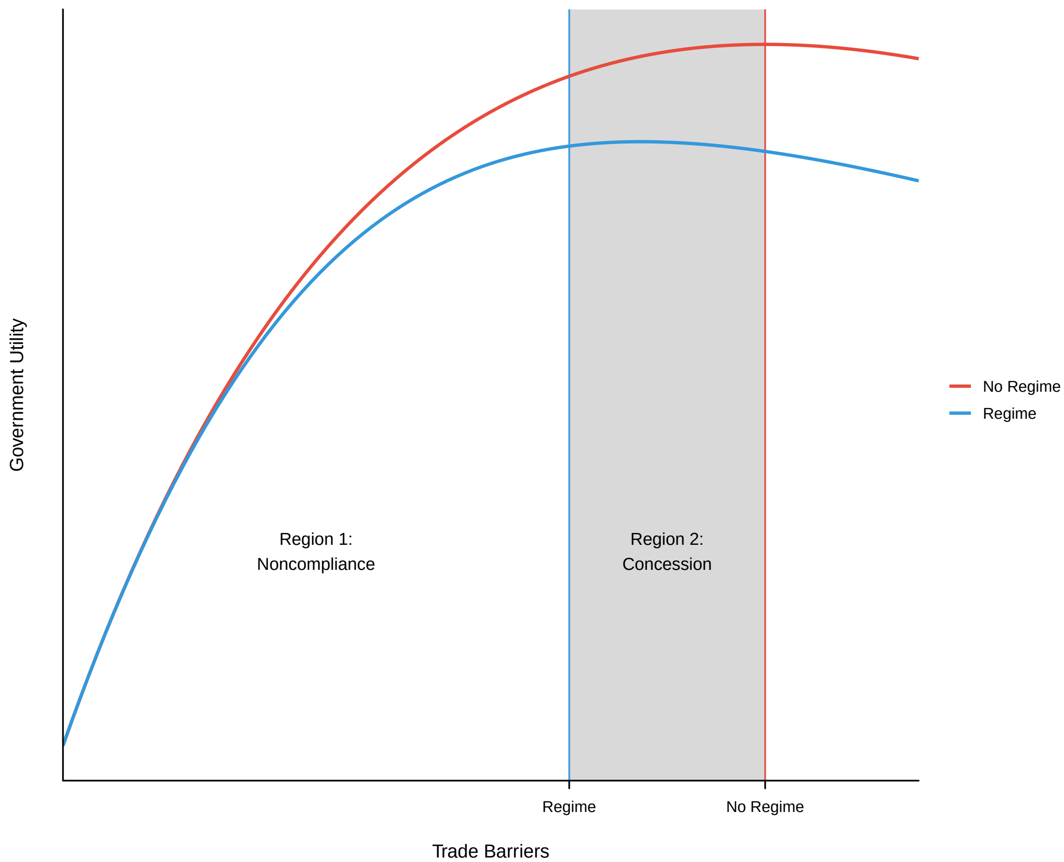

Figure 2.5. Equilibrium Concession with a Regulatory Regime

Note: Governments make a concession to avoid costly litigation. This concession is the effect of the regulatory regime on equilibrium trade barriers. The width of Region 1 indicates the amount of *ex ante* noncompliance and the width of Region 2 indicates the size of the concession.

The width of Region 1, $b_0^* - b_t$, indicates remaining *ex ante* noncompliance and the width of Region 2, $\tilde{b}^* - b_0^*$, indicates the size of the concession.

> **Result 8.** Regulatory regimes cause member state governments to make optimal concessions by reducing expected *ex ante* trade barriers: $b_0^* - \tilde{b}^* < 0$. This also reduces *ex post* trade barriers in expectation: $E\left[b_1^*\right] - \tilde{b}^* < 0$.

In equilibrium, *ex ante* trade barriers b_0^* are strictly decreasing in sector homogeneity (see Figure 2.6), just like equilibrium trade barriers $\tilde{b}^*$ in the counterfactual (see Result 7). This means that expected *ex ante* noncompliance is strictly decreasing in sector homogeneity. This is also true of *ex post* noncompliance in expectation, $E\left[b_1^*\right]$. (Note that *ex post* trade

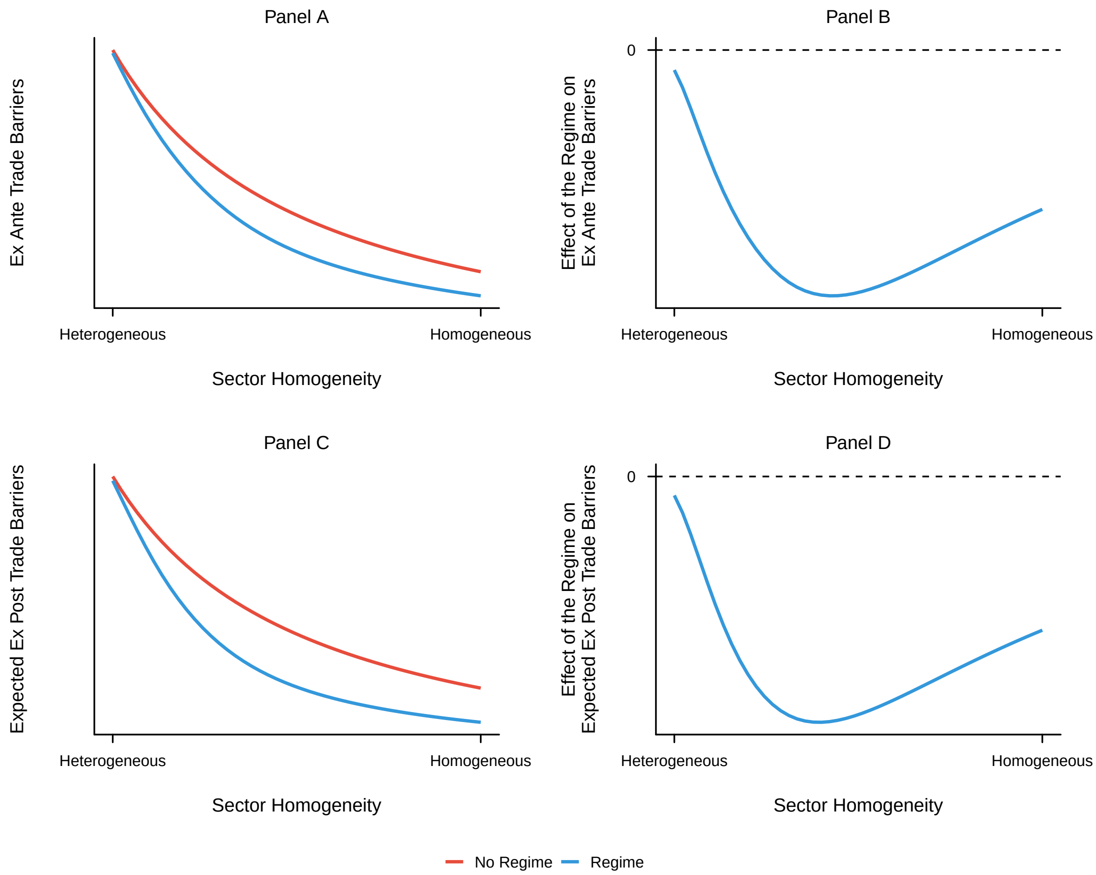

Figure 2.6. Effect of a Regulatory Regime on Equilibrium Trade Barriers

Note: The negative effect of a regime on equilibrium trade barriers (*ex ante* and *ex post*) is largest for sectors with intermediate levels of homogeneity.

barriers are strictly less than *ex ante* trade barriers in expectation, $E\left[b_1^*\right] < b_0^*$, because they are a convex combination of *ex ante* trade barriers b_0^* and the treaty-mandated trade barriers, $b_t < b_0^*$.)

However, the fact that both *ex ante* and *ex post* compliance are improving as sector homogeneity increases does not mean that regulatory regimes are most effective at reducing trade barriers in highly homogeneous sectors. In fact, a regulatory regime has the biggest impact on trade barriers for intermediate levels of sector homogeneity. Panel A of Figure 2.6 plots *ex ante* trade barriers against equilibrium trade barriers in the counterfactual. Panel C does the same for *ex post* trade barriers. The vertical difference between these lines ($b_0^* - \tilde{b}^*$ in Panel A and $E\left[b_1^*\right] - \tilde{b}^*$ in Panel C), plotted in Panels B and D, represents the effect of a regulatory regime on equilibrium trade barriers. That difference is greatest in sectors with intermediate levels of homogeneity.

Result 9. As sector homogeneity θ increases, *ex ante* noncompliance, $\left| b_0^* - b_t \right|$, decreases. The negative effect of a regulatory regime on equilibrium *ex ante* trade barriers, $b_0^* - \tilde{b}^* < 0$, is largest for sectors with intermediate levels of homogeneity. The same is true for expected *ex post* trade barriers in equilibrium, $E\left[b_1^*\right]$.

The intuition behind this result is as follows. As firms become more homogeneous, governments prefer lower trade barriers. This is because price competition is less intense, and there is less need to protect domestic firms from foreign competition (see Result 7). This allows governments to reorient their trade policies towards consumer interests. Thus, as firms become more homogeneous, the costs of compliance decrease (see Figure 2.7, Panel A). As the cost of compliance decreases, the court becomes more likely to rule against a government (more likely to rule in favor of a litigant) because it is more likely that the government will comply with the court's ruling (see Figure 2.7, Panel B). In other words, *ex post* compliance is more likely.

The fact that the court becomes more likely to rule in favor of litigants incentivizes litigants to bring noncompliance cases against governments. When firms are heterogeneous, *ex ante* noncompliance is high, so litigants would like to correct it. But they are not likely to actually bring cases because the probability that the court will rule against governments is low. Thus, bringing a case is unlikely to be worth the cost of litigating. As firms become more homogeneous, litigants become more likely to bring a case because the court is more likely to rule against governments.

At the same time, however, governments prefer to choose lower, more complaint trade barriers (see Figure 2.4, Panel A), which means that the benefits of correcting noncompliance are dropping. At some point, the costs of litigation are no longer justified, even if the court is likely to rule against governments and governments are likely to comply with that ruling.

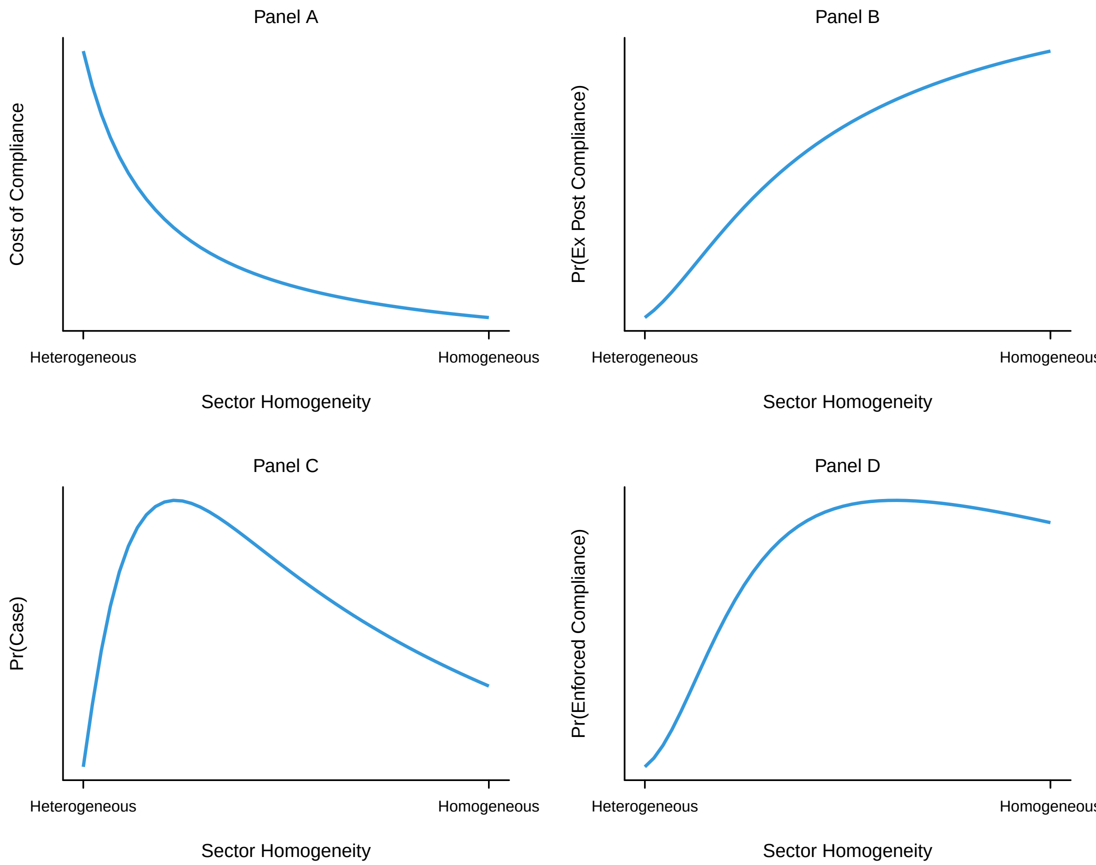

Figure 2.7. Comparative Statics with a Regulatory Regime

Note: With a regulatory regime, the costs of compliance are decreasing in the costs of compliance. The conditional probability of *ex post* compliance is increasing. The probability of a case and the probability of enforced compliance are largest for sectors with intermediate levels of homogeneity.

These competing incentives to litigate mean that the probability of a case is highest for sectors with intermediate levels of homogeneity (see Figure 2.7, Panel C).

Governments anticipate the behavior of litigants and the court. They want to avoid enforced compliance, which is the joint probability that the litigant brings a case, that the court rules against the government, and that the government comes into compliance (see Figure 2.7, Panel D). Since the probability that the litigant brings a case is highest in sectors with intermediate levels of homogeneity, so too is the probability of enforced compliance. As the probability of enforced compliance increases, governments make larger and larger concessions in an attempt to avoid litigation. Thus, the effect of the regulatory regime on *ex ante* trade barriers and expected *ex post* trade barriers is also largest for sectors with intermediate levels of homogeneity (see Figure 2.6).

In sum, the politics of compliance generate systematic bias in the types of noncompliance cases that get litigated. Litigants drop cases (a) in very heterogeneous sectors, where the costs of compliance are high, and a court is unlikely to rule against a government; and (b) in very homogeneous sectors, where the costs of litigation are high relative to the degree of noncompliance. Thus, while compliance is always better in more homogeneous sectors, regulatory regimes reduce trade barriers most in sectors with intermediate levels of homogeneity — causing a systematic distortion in the economy.

2.6.4 Distributive Consequences and Consumer Welfare Gains

Regulatory regimes affect firm performance by reducing trade barriers. Reducing trade barriers helps productive, exporting firms to gain market share at the expense of unproductive, non-exporting firms (e.g., Melitz 2003; Chaney 2008). But the politics of compliance create a systematic distortion in the economy — regulatory regimes reduce compliance most in sectors with intermediate levels of homogeneity — which has downstream effects on firm performance and consumer welfare. In this section, I identify the distributive consequences of regulatory regimes for firms and the implications for consumer welfare. Then, I show how the politics of noncompliance distorts these economic consequences of common markets by differentially reducing trade barriers across sectors.

By incentivizing governments to lower trade barriers, regulatory regimes increase the domestic production cut-point, $\varphi_d^* - \tilde{\varphi}_d^* > 0$, and decrease the exporting cut-point, $\varphi_x^* - \tilde{\varphi}_x^* < 0$. More firms can profitably export, but fewer can profitably produce for the domestic market. This change in the composition of the sector has implications for consumer welfare. Since the most unproductive firms go out of business, the average productivity of firms increases. This lowers the price index P^* (the cost of one unit of utility) and increases consumer welfare W^* (income divided by the price index).

Result 10. International regulatory regimes increase the domestic production cut-point and decrease the exporting cut-point: $\varphi_d^* - \tilde{\varphi}_d^* > 0$ and $\varphi_x^* - \tilde{\varphi}_x^* < 0$. Firms that only produce for the domestic market perform worse under the regime, and firms that export perform better. The least profitable firms exit the market and the most productive firms that produce for the domestic market start to export. Of the firms that start to export, only the most productive of these perform better under the regime. These distributive consequences improve consumer welfare, $W^* - \tilde{W}^* > 0$.

Figure 2.8 shows how a regulatory regime distorts firm performance by plotting firm performance (i.e., total expected net profit from domestic and foreign markets) in equilibrium as a function of productivity in a world with a regime and in a world without a regime (the counterfactual). The regime increases the domestic production cut-point from point A to point B and decreases the exporting cut-point from point C to point D. The shaded areas between these two sets of points represent the magnitude of these effects. Firms to the left of point E (where the profit lines intersect) perform worse because of the regime and firms to the right of point E perform better. These points divide the productivity space into six regions. How exactly a regulatory regime will impact the behavior and profitability of an individual firm depends on which of these regions it falls into.

Firms in Region 1 are so unproductive that they never enter the market, and are therefore not affected by the regime. Firms in Region 2 exit the domestic market because of the regime. They are productive enough to produce without the regime (point A), but not with the regime (point B). Firms in Region 3 produce for the domestic market either way, but they perform better without the regime. Firms in Regions 4 and 5 start to export because the existence of the regime lowers the exporting cut-point by inducing member states to make concessions (i.e., to reduce trade barriers). They are not productive enough to export without the regime (point C), but they are productive enough with the regime (point D). Firms in Region 4 perform worse even though they start exporting but firms in Region

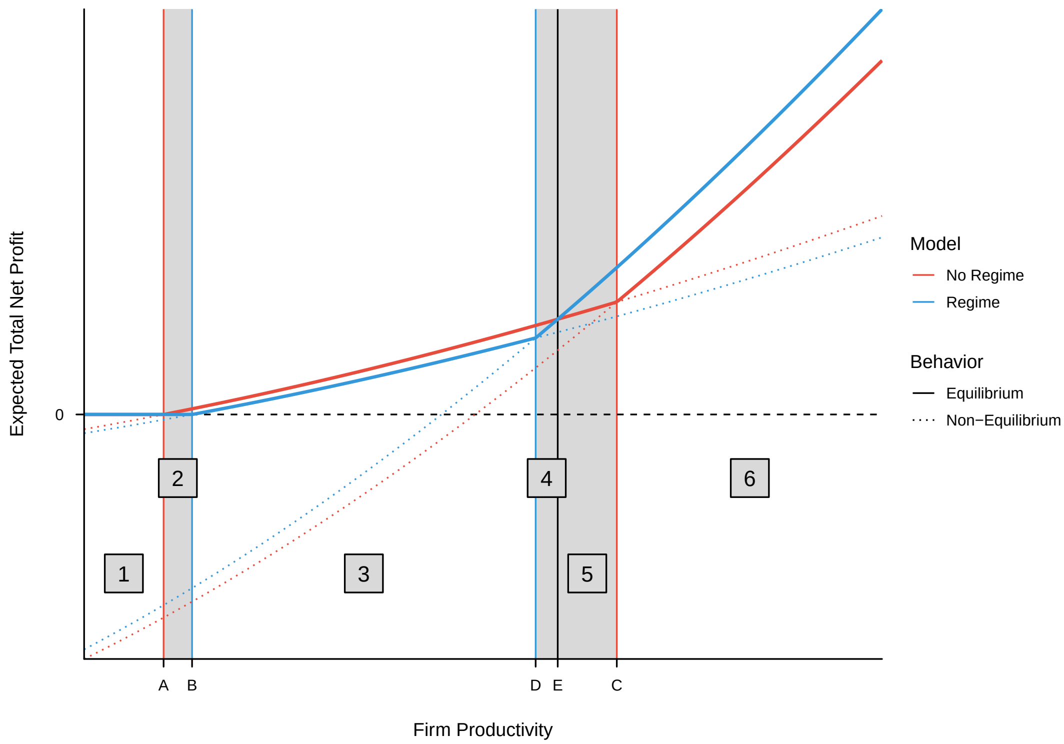

Figure 2.8. Distributive Consequences of a Regulatory Regime

Note: Firms in Regions 1, 2, 3, and 4 perform worse under an international regulatory regime, whereas firms in Regions 5 and 6 perform better. Firms in Region 1 never enter the domestic market. Firms in Region 2 exit the domestic market because of the regime. Firms in Region 3 lose market share. Firms in Regions 4 and 5 start to export. Firms in Region 5 become more profitable, but firms in Region 4 become less profitable. Firms in Region 6 continue to export but gain market share.

5 perform better because they start exporting. Firms in Region 6 export either way but perform better because of the regime.

To determine how the politics of compliance in regulatory regimes systematically distorts the distributive consequences of liberalization, I calculate the effect of a regulatory regime on the productivity cut-points — the difference in the cut-points with the regime and without the regime — as a function of sector homogeneity (see Figure 2.9, Panels A and C). This is sufficient because a firm's productivity relative to these changing cut-points fully determines the distributive consequences of liberalization for that firm — whether the firm is a winner or loser from the development of a common market.

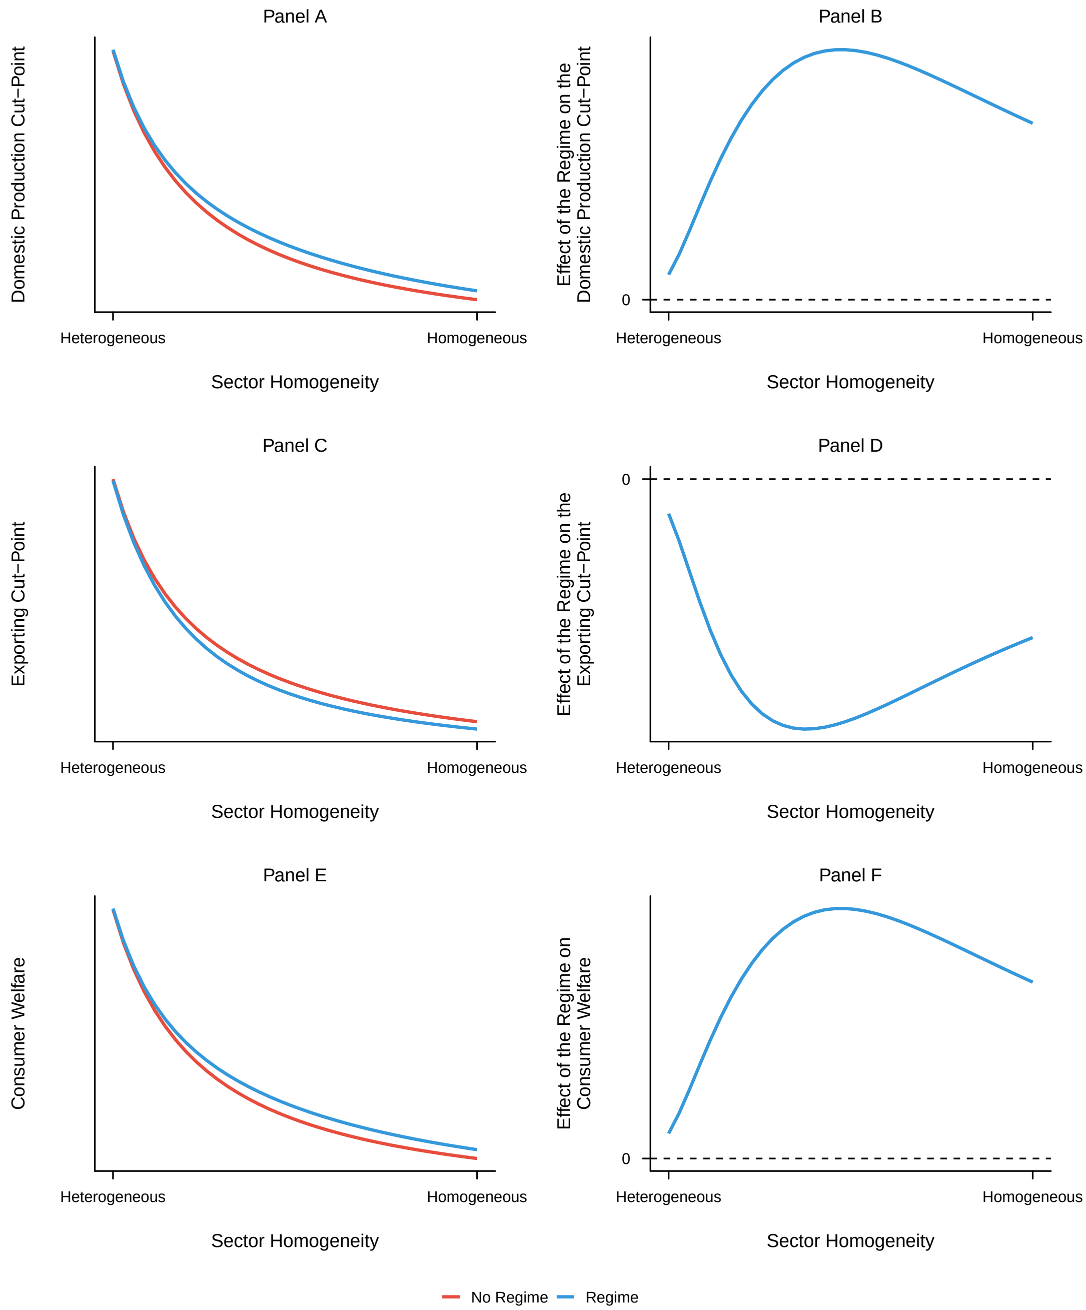

Figure 2.9. Effect of a Regulatory Regime on Firm Behavior and Consumer Welfare

Note: The effect of a regulatory regime on the domestic production cut-point (positive), the exporting cut-point (negative), and consumer welfare (positive) are largest for sectors with intermediate levels of homogeneity.

The effect of the regime on the domestic production cut-point is always positive and the effect of the regime on the exporting cut-point is always negative (see Result 10), but sector homogeneity conditions the magnitude of these effects. The positive effect of the regime on the domestic production cut-point is largest for intermediate levels of homogeneity (see Figure 2.9, Panel B). The negative effect of the regime on the exporting cut-point is also largest for intermediate levels of homogeneity (see Figure 2.9, Panel D). These distortions — due to bias in which cases get litigated — matter to individual firms. Firms in highly homogeneous sectors or highly heterogeneous sectors do not face the distributive consequences of common markets to the same degree as those in sectors with intermediate levels of homogeneity (for better or worse). Sufficiently productive firms miss out on gains, whereas sufficiently unproductive firms avoid loses.

These effects carry over to consumer welfare: the positive effect of the regime on consumer welfare is also largest for sectors with intermediate levels of homogeneity (see Figure 2.9, Panels E and F). Consumers who tend to purchase goods from highly heterogeneous sectors or highly homogeneous sectors miss out on the welfare gains from common markets, making them relatively worse off than consumers who tend to purchase goods from sectors with intermediate levels of homogeneity.

Result 11. The positive effect of a regulatory regime on the domestic production cut-point, $\varphi_d^* - \tilde{\varphi}_d^* > 0$, the negative effect of the regime on the exporting cut-point, $\varphi_x^* - \tilde{\varphi}_x^* < 0$, and the positive effect of the regime on consumer welfare, $W^* - \tilde{W}^* > 0$, are largest for sectors with intermediate levels of homogeneity θ.

In sum, by reducing trade barriers, regulatory regimes create distributive consequences: they allow new firms to export and push unproductive firms out of business. This raises average firm productivity, lowers average prices, and improves consumer welfare. But the politics of noncompliance generates systematic bias in the types of cases that get litigated, causing the regime to constrain trade barriers more in sectors with intermediate levels of homogeneity. Thus, these effects on firm performance and consumer welfare are largest in sectors with intermediate levels of homogeneity.

2.7 Conclusion

Countries create common markets to accrue consumer welfare gains. To ensure that they actually realize these gains, they rationally design international regulatory regimes to manage noncompliance with the rules of the common market. But the politics of compliance generates systematic bias in the types of noncompliance cases that get litigated. I develop a formal model that explains how the politics of compliance in regulatory regimes systematically distorts the welfare gains that states accrue — the very reason they create common markets in the first place. I show that if we do not take into account the politics of compliance, our theoretical predictions about the distributive consequences of trade liberalization — which firms win and lose — and the conditions under which member states will accrue consumer welfare gains by creating common markets will be biased.

The model predicts the sectors in which regulatory regimes will be effective at reducing trade barriers — those with intermediate levels of homogeneity in terms of firm productivity — as well as the downstream consequences for the performance of individual firms and for consumer welfare gains. Regulatory regimes are most effective at reducing trade barriers in sectors with intermediate levels of firm homogeneity. In highly homogeneous sectors and highly heterogeneous sectors, regulatory regimes are not effective at helping member states accrue consumer welfare gains.

However, the reason why regulatory regimes are ineffective in highly heterogeneous sectors is very different than the reason they are ineffective in highly homogeneous sectors. In heterogeneous sectors, price competition is higher, giving governments more incentive to erect trade barriers to protect domestic firms. Regulatory regimes are ineffective because governments are unlikely to comply with adverse court rulings, making the court hesitant to rule against them (Carrubba and Gabel 2015; Fjelstul and Carrubba 2018). This deters litigants from bringing noncompliance cases, creating a compliance deficit (König and Mäder 2014; Fjelstul and Carrubba 2018).

In homogeneous sectors, on the other hand, regulatory regimes are ineffective because the benefits of successfully prosecuting noncompliance are low relatively to the costs of litigation. In homogeneous sectors, price competition is low and governments are more willing to comply. The court is more likely to rule in favor of the litigant, but compliance is good enough that the benefits of bringing governments into compliance are not worth the costs of litigating. In homogeneous sectors, the regime is ineffective because it is not used, not because it cannot correct violations.

The fact that common markets are less ineffective at generating welfare gains in certain sectors of the economy may put political pressure (from firms and consumers) on member states whose economies depend on those sectors — in terms of production or in terms of consumption. This may help explain variation in public support for common markets across countries. The long-term consequences of uneven welfare gains for the political stability of common markets is an important question for future work.

In sum, if member states always complied with the rules of a common market, variation in sector homogeneity would not create systematic distortions in which firms win and lose from trade liberalization through the development of a common market, and consumer welfare gains would be constant across sectors of the economy. Thus, taking into account the politics of compliance in regulatory regimes is critical to our understanding of how well common markets work. This has implications for the long-term stability of several of the world's largest economies.

Paper 3:

Rationality in the Public's Evaluation of Economic Policy: Evidence from the European Sovereign Debt Crisis

Abstract

The increased use of direct democracy to make international economic policy raises an important question: To what extent are aggregate preferences over complex economic policies consistent with individual rationality? This question has far-reaching implications for the coherence of economic policy. In this paper, I theorize how a rational individual would evaluate the following policy: Commission monitoring of member state compliance with the Stability and Growth Pact (SGP), which governs the European Economic and Monetary Union (EMU). Rational individuals would (a) identify the likely outcomes that a policy can generate, (b) understand how the current economic situation in their country conditions the effects of a policy, (c) evaluate any tradeoffs that a policy creates, and (d) support policies that advance their self-defined interests. I present empirical evidence from the European Sovereign Debt Crisis that, in a high-information environment, individuals express preferences over Commission monitoring that, in the aggregate, are consistent with individual rationality.

3.1 Introduction

An important trend in global governance is the increasing use of direct democracy to set international economic policy. Recent examples include the 2007 referendum in Costa Rica on joining the Dominican Republic-Central American Free Trade Agreement (CAFTA-DR) (Hicks, Milner and Tingley 2014; Urbatsch 2013), the 2010 referendum in Iceland on loan guarantees (Curtis, Jupille and Leblang 2014), the 2012 referendum in Croatia on membership in the European Union (EU), and the 2016 referendum in the United Kingdom on EU membership. Beyond formal referendums, general elections often become quasi-referendums on major economic policies. For example, the 2016 election in the United States was, at least in the industrial Midwest, a referendum on decades of American trade policy. The 2017 election in the United Kingdom was, in large part, a referendum on the Conservative Party's hard-line stance on Brexit.

The use of direct democracy to set economic policy raises an important question: To what extent are aggregate preferences over complex economic policies consistent with individual rationality, even if we do not believe that individuals have the interest, ability, or information to rationally evaluate policies? This question has far-reaching implications for the coherence of economic policy. In this paper, I theorize how a rational individual would evaluate the following policy: Commission monitoring of member state compliance with the Stability and Growth Pact (SGP) — an agreement between European Union (EU) member states that governs the European Economic and Monetary Union (EMU). I present empirical evidence from the European Sovereign Debt Crisis that, in a high-information environment, individuals express preferences over Commission monitoring that, in the aggregate, are consistent with individual rationality.[1]

[1] Note that this evidence *does not* suggest than individuals cognitively evaluate complex economic policies rationally, only that aggregate preference patterns are consistent with individual rationality.

Complex economic policies are multi-dimensional: they can produce multiple outcomes that individuals can have preferences over (Carrubba and Singh 2004). They can also generate different political outcomes in different contexts. National economic conditions, for example, can shape the real-world political outcomes of a policy. Depending on individuals' preferences over those outcomes, a policy can present a tradeoff under some conditions, but not others. I argue that rational individuals would (a) identify the likely outcomes that a policy can generate, (b) understand how the current economic situation in their country conditions the effects of a policy, (c) evaluate any tradeoffs that a policy creates, and (d) support policies that advance their self-defined interests.

The usual approach to assessing the rationality of policy preferences, in contrast, is to ascribe interests to individuals based on their observable characteristics and then to use those characteristics to predict self-reported support for policies that advance them. For example, an individual with credit card debt might be more likely to support sovereign debt repayment to keep their interest rate low (Curtis, Jupille and Leblang 2014). Scholars then evaluate the extent to which policy preferences are driven by material self-interest versus sociotropic considerations.[2] A rational individual, according to this approach, is one that can correctly identify their material self-interests (according to standard economic theory) and support policies that advance those interests.[3]

[2] The empirical evidence is mixed. Some studies find empirical evidence that policy preferences are driven by material self-interest (Gabel 2000; Mayda and Rodrik 2005; Hays, Ehrlich and Peinhardt 2005; Mayda 2006; O'Rourke and Taylor 2006; Jupille and Leblang 2007; Mayda 2008; Pandya 2010; Fordham and Kleinberg 2012; Curtis, Jupille and Leblang 2014; Bearce and Tuxhorn 2017), whereas others find evidence that preferences are driven by sociotropic factors (Hiscox 2006; Hainmueller and Hiscox 2006, 2007; Mansfield and Mutz 2009; Hainmueller and Hiscox 2010). One of the original motivations behind this research agenda was to test the micro-foundations for the Heckscher-Ohlin and Ricardo-Viner models (Scheve and Slaughter 2001), but there is little empirical evidence that economic policy preferences are driven by ownership of relatively abundant factors of production (Heckscher-Ohlin) or employment in export-oriented versus import-competing sectors (Ricardo-Viner).

[3] According to the literature, individuals are more likely to correctly identify their material self-interests when they are well-informed (Gomez and Wilson 2006; Armingeon and Ceka 2014; Bearce and Tuxhorn 2017), such as when they are operating in a high-information environment, like an international financial crisis (Curtis, Jupille and Leblang 2014).

In this paper, I evaluate the degree to which aggregate policy preferences over Commission monitoring of member state compliance with the SGP are consistent with individual rationality. The SGP consists of two criteria, which place legal limits on member states' deficit spending and sovereign debt. These criteria help mitigate a perverse incentive that member states have to over-spend and over-borrow in order to stimulate their economies. Member states know that, in the event of a sovereign debt crisis, the EU would have no choice but to bail them out because each member state is too big to fail — a textbook example of moral hazard. I choose this substantive application because the Eurozone crisis is a high-information environment where the costs and benefits of proposed economic policies are most tangible (Curtis, Jupille and Leblang 2014).

The political consequences of Commission monitoring, including the risk of a sovereign default, depend on a member state's noncompliance with the SGP criteria. An individual's self-reported support for the EMU (an observable indicator of her interest in the long-term stability of the EMU) will shape whether they view these outcomes as costs or benefits, and therefore the probability that they will express support for monitoring.

Public opinion about Commission monitoring is substantively important because of the implications for member state compliance with the SGP. When the public does not support monitoring, governments have a greater incentive to violate the SGP criteria. As the cost of compliance increases for member states, the Commission will become less likely to take legal action against noncompliant member states. Knowing that it will be more difficult to successfully bring member states into compliance, the Commission will become less wiling to pay the financial and political costs of monitoring and enforcement activities (König and Mäder 2014; Fjelstul and Carrubba 2018; Baerg and Hallerberg 2016). This strategic behavior can cause a compliance deficit.

Using multi-level models, I provide empirical evidence that the aggregate policy preferences of survey respondents are consistent with individual rationality. Specifically, I show that an individual's support for Commission monitoring depends on the interaction between their self-reported support for the EMU and her member state's compliance with the SGP criteria. In developing predictions, I black-box the origin of individuals' self-defined interests. However, public discourse about policy responses to the Eurozone crisis could affect whether an individual supports the EMU (i.e., how they defines their interests). To address this potential endogeneity problem, I use propensity score matching to correct for the non-random assignment of support for the EMU.

3.2 The Substantive Application

To theorize individuals' support for Commission monitoring as a potential solution to the moral hazard problem underlying the Eurozone crisis, we need to consider three key issues: (1) the role moral hazard has played in the crisis, (2) how Commission monitoring mitigates moral hazard, and (3) the tradeoffs that Commission monitoring presents.

First, what role has moral hazard played in the sovereign debt crisis? The stability of the EMU depends on a non-credible promise: That member states will adhere to sound fiscal policies when the economy is good and credit plentiful. If a member state violates this promise, and its sovereign debt becomes so unwieldy that investors doubt that it will be able to pay the interest on its debt, they could launch a speculative attack — a sudden, uncoordinated sell-off of government bonds in anticipation of a precipitous drop in their value. By flooding the bond market, investors increase supply, which decreases the price. Investors then demand higher interest rates to compensate for the risk of owning a bond that could become next to worthless in the event of a default. Spiking interest rates make it that much harder for the government to restructure its sovereign debt.

This is exactly what happened to Greece on the heels of the Great Recession. Investors launched a speculative attack on Greek bonds when a newly-elected government announced that Greece's fiscal position had been vastly overstated. (Its deficit and debt were much higher than reported.) Greece was in a bind. Usually, debt-burdened countries decrease the real value of their debt burden by devaluing their currencies. As a member of the EMU, Greece did not have that option. It could not unilaterally use monetary policy to pull itself out of the hole. Greece's inability to borrow as interest rates on its bonds spiked pushed the government towards sovereign default and sparked the European sovereign debt crisis, which threatens to destabilize the entire EMU.

A Greek default would not be an isolated event. Economists and politicians alike have warned that letting Greece default and exit the EMU (the only way to reclaim monetary policy autonomy in order to reduce its real debt burden) could create financial contagion. A Greek sovereign default would hurt investors who hold Greek sovereign bonds — investors that include other EU member states. Non-performing Greek assets would imperil investors' already-weak financial positions. Beyond financial contagion, a so-called *Grexit* would create a dangerous precedent. It would prove that EMU membership is reversible.

The dramatic steps that the Troika (a consortium comprising the International Monetary Fund, the European Central Bank, and the European Commission) has taken to bailout Greece indicate the wide-spread belief that at-risk member states cannot be allowed to default and exit the EMU — that each EMU member is too big to fail.

The expectation of a bailout in the event of sovereign default perversely incentivizes over-spending and over-borrowing in the short term to stimulate the domestic economy. This is a textbook example of *moral hazard*. Knowing that a bailout is likely — because a sovereign default could create enough financial contagion to bring the European economy to its knees — incentivizes EMU members to over-spend. In short, EMU members do not internalize the costs of their own risky behavior.

Second, how can Commission monitoring mitigate moral hazard? The architects of the EMU created the Stability and Growth Pact (SGP) — an agreement among all EU member states designed to underpin the stability of the EMU by placing caps on member states' deficits (3 percent of GDP) and sovereign debts (60 percent of GDP) — to mitigate moral hazard. The EU treaties task the European Commission with monitoring compliance. Monitoring is the primary tool the EU has to mitigate moral hazard. The SGP allows the Commission to fine noncompliant member states to incentivize compliance.

The SGP has been spectacularly ineffective, however, because the Commission has declined to initiate noncompliance proceedings to enforce the SGP criteria. Since *ex post* monitoring has failed to dis-incentivize noncompliance, observers have called for reforms that would empower the Commission to pre-approve member state budgets to prevent noncompliance before it happens (*ex ante* monitoring).[4]

Third, what tradeoffs does Commission monitoring of member state compliance (*ex post* or *ex ante*) with the SGP criteria present to individuals evaluating monitoring as a policy to address to moral hazard? Effective monitoring (and enforcement) by the Commission constrains a country's deficit spending, thereby limiting marginal increases to its sovereign debt, which helps to manage the risk of sovereign default. In short, effective monitoring means Commission-imposed austerity.[5] Austerity is costly, however. By contracting public spending, austerity hurts short-term economic growth. This stability/austerity tradeoff is at the heart of the public discourse about austerity measures across Europe.

3.3 Theory

How would a rational individual evaluate the policy of Commission monitoring of the SGP criteria? This section builds on the IPE literatures on economic crises and public opinion to theorize when a rational individual would be more likely to support two forms of Commission monitoring: (1) pre-approval of member states' budgets and (2) financial sanctions for noncompliance with SGP deficit criterion. I will refer to these as *ex ante* and *ex post* monitoring, respectively. Note that financial sanctions for noncompliance are an important part

[4] Financial sanctions for noncompliance constitutes the so-called *dissuasive arm* of the SGP. The other arm is the *preventive arm.* The Commission can help facilitate fiscal coordination between member states, but does not have the power to pre-approve member state budgets.

[5] The Commission can use *ex post* financial sanctions for noncompliance with the SGP to coerce member states into adopting austerity measures to reduce their structural deficits. It could use *ex ante* pre-approval powers for the same purpose.

of the current SGP rules. Pre-approval of member states' budgets, in contrast, is a policy proposal — an alternative method of enforcing the SGP criteria, given that the Commission has historically opted not to pursue financial sanctions.

Note that *ex ante* and *ex post* monitoring incentivize member states to adopt austerity measures using different approaches. Austerity measures can include tax hikes or spending cuts. With *ex ante* monitoring, the Commission imposes financial sanctions to punish EMU members that do not take corrective action by implementing austerity measures. If the Commission could pre-approve member state budgets, it could condition budget approval on a member state's adoption of sufficiently severe austerity measures. Generally speaking, *ex ante* monitoring provides tighter control because *ex post* monitoring (i.e., threats to impose punishments) can be non-credible.

3.3.1 Independent Variables

A rational individual's support for Commission monitoring will depend on three variables: (1) a member state's noncompliance with the SGP deficit criterion, (2) a member state's noncompliance with the SGP debt criterion, and (3) whether an individual supports the EMU (i.e., has a self-reported interest in maintaining the stability of the EMU). A member state's compliance with the SGP criteria matters because it drives the expected political outcomes of monitoring (discussed below). An individual's support for the EMU affects whether they view these outcomes as costs or benefits.

A member state's noncompliance with the SGP deficit criterion determines whether the Commission can legally initiate enforcement actions that incentivize member states to adopt austerity measures, which increase the probability of a recession; the Commission can only do so when a member state's deficit is SGP-noncompliance. Thus, monitoring only has the potential to be costly, in terms of Commission-imposed austerity, if a member state is not in compliance with the SGP deficit criterion (Column 4 of Table 3.1).

It is important to note that the Commission cannot enforce the SGP debt criterion directly; rather, it enforces the debt criterion indirectly by enforcing the deficit criterion. Once a member state has accumulated a large public debt, the Commission cannot do anything about it short of helping the member state negotiate with investors to restructure their sovereign debt. The objective of monitoring is to limit new contributions to member states' existing debts by constraining deficit spending. Every time a member state runs a deficit, it must borrow to cover outlays in excess of revenue by issuing government bonds, which directly contributes to the sovereign debt.

A member state's noncompliance with the debt criteria affects the probability of sovereign default. There is a risk of sovereign default when a member state is violating only the debt criteria, as any deficit is a marginal contribution to the debt. But that risk is higher when the member state is also violating the deficit criteria, as the marginal contribution to the debt is relatively larger (Column 5 of Table 3.1). When a member state's sovereign debt is expanding at a sufficiently accelerating rate, a speculative attack by investors becomes more likely. Raising interest rates can then push a member state to default.

By affecting (a) whether the Commission can impose austerity and (b) the risk of sovereign default, a member state's noncompliance with the SGP criteria drives the expected effect of Commission monitoring on three political outcomes that individuals will have preferences over: (a) the probability of austerity, (b) the probability that an individual member state will exit the EMU, and (c) the probability that the EMU will break up (i.e., the long-term stability of the EMU). As I discuss below, whether individuals view these political outcomes as costs or benefits, which depends on their support for the EMU, will determine their support for the policy of Commission monitoring.

First, Commission monitoring increases the probability of Commission-imposed austerity, but only when a member state's debt and deficit as SGP-noncompliant. The Commission has the ability to impose austerity when a member state's deficit is non-compliant, but recent research shows that the Commission selectively enforces compliance with EU law

Table 3.1. Summary of Theory

Interests	Domestic economic context		Implications		Expected effect of monitoring		
(1) EMU	(2) Deficit	(3) Debt	(4) Austerity	(5) Pr(Default)	(6) Pr(Austerity)	(7) Pr(Exit)	(8) Pr(Breakup)
Support	Noncompliant	Noncompliant	Yes	High	Increase (−)	Decrease (+)	Decrease (+)
Support	Compliant	Noncompliant	No	Low	Constant	Decrease (+)	Decrease (+)
Support	Noncompliant	Compliant	Yes	None	Constant	Constant	Decrease (+)
Support	Compliant	Compliant	No	None	Constant	Constant	Decrease (+)
Oppose	Noncompliant	Noncompliant	Yes	High	Increase (−)	Decrease	Decrease
Oppose	Compliant	Noncompliant	No	Low	Constant	Decrease	Decrease
Oppose	Noncompliant	Compliant	Yes	None	Constant	Constant	Decrease
Oppose	Compliant	Compliant	No	None	Constant	Constant	Decrease

Note: Columns 1–3 indicate the independent variables of interest: an individual's support for the EMU, their member state's compliance with the SGP deficit criteria, and their member state's compliance with the SGP debt criteria. Each row corresponds to a unique combination of values. Column 4 indicates whether the Commission can incentivize austerity due to a violation of one of the SGP criteria. Column 5 indicates the probability of sovereign default. Columns 6–8 indicate the expected effect of Commission monitoring on the probability of Commission-imposed austerity, the probability the member state will need to exit the EMU in order to recover from a sovereign default, and the probability the EMU will break up due to financial contagion caused by a sovereign default (by any member state). The symbols in parentheses indicate whether an individual, based on their support for the EMU, is likely to view each these effects as a cost (−), a benefit (+), or whether they would be indifferent (no symbol).

(König and Mäder 2014; Baerg and Hallerberg 2016; Fjelstul and Carrubba 2018). The Commission is more likely to actually enforce the SGP deficit criterion when the member state faces legitimate risk of a sovereign default, which is when a member state's debt is also SGP-noncompliant (Column 6 of Table 3.1).

Second, Commission monitoring decreases the probability that a member state will be forced to exit the EMU when a member state's debt is SGP-noncompliant. When a member state's deficit is also SGP-noncompliant, monitoring reduces the risk of a sovereign default, which could force a member state to exit the EMU in order to depreciate its domestic currency, thereby reducing its real debt burden. When a member state's deficit is SGP-compliant, on the other hand, monitoring promotes fiscal discipline in the future. Governments have an incentive to over-spend due to a time inconsistency problem. Monitoring creates a disincentive for the member state to over-spend to counterbalance this incentive. Note that monitoring presents a tradeoff for individuals in EMU members whose deficits and debts are SGP-noncompliant: a lower risk of default and exit from the EMU in the long term in exchange for a higher risk of economic recession in the short term (Columns 6 and 7 of Table 3.1). This is the stability/austerity tradeoff mentioned above.

Third, Commission monitoring decreases the probability of the EMU breaking up by reducing moral hazard. This is true regardless of a member state's noncompliance with the SGP criteria (Column 8 of Table 3.1). A sovereign default, even by a small member state, could threaten the stability of the entire EMU because of the potential for financial contagion. By reducing the risk of sovereign default, monitoring helps to prevent a situation in which financial contagion could spread to other member states. A sovereign default would force other member states and their financial institutions to take a loss non-performing sovereign bonds, which could also cause a banking crisis. Financial contagion could prompt the worst-hit member states to withdraw from the EMU in order to re-introduce their own domestic currencies, which they could then depreciate to reduce their real debt burden.

An individual's support for the EMU will condition whether they view each of these three political outcomes of Commission monitoring as costs or benefits. The exception is that an individual's support for the EMU will not condition how they view an increase in the probability of Commission-imposed austerity; all individuals should view an increase in this probability as a cost of monitoring (Column 6 of Table 3.1). The downside of austerity is that it increases the risk of a severe economic recession in the short term. From a Keynesian perspective, the more severe a member state's austerity policies, the bigger the contractionary impact of reduced government spending on its domestic economy. This cost applies when a member state's debt and deficit are both SGP-noncompliant, as that is when the Commission is likely to enforce the SGP criteria.

That being said, an individual who supports the EMU will be more likely to see a decrease in the probability that their member state will be forced to exit the EMU as a benefit of monitoring. Monitoring reduces the likelihood that a member state will need to exit the EMU in order to recover from a sovereign default by re-introducing and depreciating its own currency. This benefit applies when a member state's debt is SGP-noncompliant (Column 7 of Table 3.1). Individuals who oppose the EMU, on the other hand, are more likely to be indifferent to a change in this probability.

A similar logic applies to how individuals will view a decrease in the probability of the EMU breaking up due to contagion caused by the sovereign default of another EMU member state. By addressing the EMU's moral hazard problem, Commission monitoring decreases the probability that a sovereign default by any EMU member state will spread financial contagion to other member states. This is true regardless of a member state's compliance with the SGP criteria. However, individuals who support the EMU are more likely to see a decrease in this probability as a benefit, whereas individuals who oppose the EMU are more likely to be indifferent (Column 8 of Table 3.1).

3.3.2 Hypotheses

I develop testable hypotheses about how an individual's support for the EMU will interact with their member state's noncompliance with the SGP criteria to shape their support for Commission monitoring. Since interactions are symmetric, I develop predictions for how the effect of each of these three variables on support for monitoring is moderated by the other two. This ensures that I do not ignore testable predictions (Brambor, Clark and Golder 2006; Franzese and Kam 2009; Berry, Golder and Milton 2012).

First, I consider how the effect of an increase in a member state's noncompliance with the SGP debt criterion on an individual's support for Commission monitoring depends on (a) whether the individual supports the EMU and (b) whether their member state's deficit is SGP-noncompliant (see Hypothesis 1 in Table 3.2).[6]

When an individual supports the EMU and their member state's deficit is SGP-noncompliant, the effect of an increase in their member state's noncompliance with the SGP debt criterion on their support for monitoring is ambiguous. In this case, the individual faces the stability/austerity tradeoff: an increase in her member state's noncompliance with the SGP debt criterion reduces the probability that a sovereign default will force their member state to exit the EMU (see Column 7 of 3.1), but their member state is also exposed to Commission-imposed austerity (see Column 6 of 3.1). The net effect is ambiguous.

When an individual supports the EMU and their member states's deficit is SGP-compliant, an increase in their member state's noncompliance with the SGP debt criterion increase their support for monitoring. In this case, there is no tradeoff. Their member state is not violating the SGP deficit criterion, so there is no real risk of Commission-imposed austerity (see Column 6 in Table 3.1). Strict monitoring of the SGP debt criteria locks in current deficit spending levels, limiting the size of accumulating deficits. This reduces the probability that

[6] Note that a change in a member state's compliance with the SGP debt criterion, conditional on their support for the EMU, will not change the effect of monitoring on the probability of the EMU breaking up or how an individual views that outcome (see Column 8 in Table 3.1).

a sovereign default will force their member state to exit the EMU, which they will view as a benefit (see Column 7 in Table 3.1).

When an individual does not support the EMU and their member state's deficit is SGP-noncompliant, an increase in their member state's noncompliance with the SGP debt criterion will decrease their support for monitoring. Since the member state's deficit is SGP-noncompliant, an increase in their member state's noncompliance with the SGP debt criterion increases the risk of Commission-imposed austerity (see Column 6 in Table 3.1). This will also decrease the probability that a sovereign default will force their member state to exit the EMU (see Column 7 in Table 3.1). However, since the individual does not support the EMU, they will be indifferent to a change in this probability. Thus, their support for monitoring will decrease.

When an individual opposes the EMU and their member state's deficit is SGP-compliant, an increase in their member state's noncompliance with the SGP debt criterion will have no effect on their support for monitoring. Because their member state's deficit is SGP-compliant, the risk of Commission-imposed austerity remains constant as her member state's noncompliance with the SGP debt criterion increases (see Column 6 in Table 3.1). The probability that a sovereign debt crisis could force their member state to exit the EMU decreases, but because the individual does not support the EMU, they are indifferent to this effect (see Column 7 in Table 3.1).

Hypothesis 1a. If an individual supports the EMU and their member state's deficit is SGP-compliant, an increase in noncompliance with the SGP debt criterion will increase their support for monitoring.

Hypothesis 1b. If an individual opposes the EMU and their member state's deficit is SGP-noncompliant, an increase in noncompliance with the SGP debt criterion will decrease their support for monitoring.

Second, I consider how the effect of an individual's support for the EMU on their support for Commission monitoring depends on (a) whether their member states's deficit is

Table 3.2. Summary of Hypotheses

Hypothesis 1			
(1) Debt	(2) EMU	(3) Deficit	(4) Support for monitoring
Compliant → Compliant	Oppose	Compliant	Constant
Compliant → Noncompliant	Oppose	Noncompliant	Decrease
Compliant → Noncompliant	Support	Compliant	Increase
Compliant → Noncompliant	Support	Noncompliant	Ambiguous
Hypothesis 2			
(1) EMU	(2) Deficit	(3) Debt	(4) Support for monitoring
Oppose → Support	Complaint	Compliant	Weak increase
Oppose → Support	Compliant	Noncompliant	Strong increase
Oppose → Support	Noncompliant	Compliant	Weak increase
Oppose → Support	Noncompliant	Noncompliant	Moderate increase
Hypothesis 3			
(1) Deficit	(2) EMU	(3) Debt	(4) Support for monitoring
Compliant → Noncompliant	Oppose	Compliant	Constant
Compliant → Noncompliant	Oppose	Noncompliant	Decrease
Compliant → Noncompliant	Support	Compliant	Constant
Compliant → Noncompliant	Support	Noncompliant	Decrease

Note: Column 1 indicates a change in one of the independent variables. Columns 2 and 3 indicate values of the other two independent variables (i.e., the moderating variables). Column 4 indicates how support for monitoring should change in response to the change indicated in the first column, conditional on the two moderating variables.

SGP-noncompliant and (b) whether their member state's debt is SGP-noncompliant (see Hypothesis 2 in Table 3.2). Commission monitoring increases the stability of the EMU by addressing the EMU's moral hazard problem (see Column 8 of Table 3.1); thus, individuals who support the EMU will be more likely to support monitoring, regardless of their member state's compliance with the SGP criteria. However, their member state's compliance with the SGP criteria will affect the magnitude of this positive effect.

Individuals who support the EMU will gain an additional benefit from monitoring when their member state's debt is SGP-noncompliant, as monitoring also reduces the probability that their member state will be forced to exit the EMU in response to a sovereign debt crisis (see Column 7 of Table 3.1). Thus, the positive marginal effect of an individual's support for the EMU will be larger in magnitude when their member state's debt is SGP-noncompliant than when their member state's debt is SGP-compliant.

This interaction effect is conditional on whether a member state's deficit is SGP-compliant. When a member state's deficit is SGP-noncompliant, the increased benefits of monitoring that come with a larger, more noncompliant debt (i.e., a decrease in the probability that a member state will be forced to exit the EMU in response to a sovereign debt crisis) will be partially offset by a higher probability of Commission-imposed austerity (see Column 6 of Table 3.1). Thus, conditional on noncompliance with the SGP debt criterion, the positive marginal effect of an individual's support for the EMU will be larger in magnitude when their member state's deficit is SGP-compliant (i.e., when there is no risk of Commission-imposed austerity) than when the member state's deficit is SGP-noncompliant.

Hypothesis 2a. Individuals who support the EMU will be more likely to support Commission monitoring, regardless of their member state's compliance with the SGP criteria, than individuals who oppose the EMU.

Hypothesis 2b. The positive effect of an increase in an individual's support for the EMU on their support for Commission monitoring will be smallest when their member state's debt is SGP-compliant (regardless of their member state's compliance with the

SGP deficit criterion) and largest when their member state's deficit is SGP-complaint and their member state's debt is SGP-noncompliant.

Third, I consider how the effect of a member state's noncompliance with the SGP deficit criterion on an individual's support for Commission monitoring depends on (a) that member state's noncompliance with the SGP debt criterion and (b) the individual's support for the EMU (see Hypothesis 3 in Table 3.2). In this case, a member state's noncompliance with the SGP debt criterion drives the magnitude of the effect.

When the member state's debt is SGP-compliant, an increase in the member state's noncompliance with the SGP deficit criterion will not affect an individual's support for monitoring (see Rows 1 and 3 under Hypotheses 3 in Table 3.2). If an individual opposes the EMU, they will be indifferent to Commission monitoring because an increase in the member state's noncompliance with the SGP deficit criteria will not increase the risk of Commission-imposed austerity (because the member state's debt is compliant). If an individual supports the EMU, on the other hand, they will support monitoring because, while it does not affect their member state, it does address moral hazard, which increases the stability of the EMU (and because they support the EMU, they view that as a benefit).

When the member state's debt is SGP-noncompliant, an increase in the member state's noncompliance with the SGP deficit criterion will decrease support for monitoring (see Rows 2 and 4 under Hypotheses 3 in Table 3.2). This is because Commission-imposed austerity is more likely. Because the member state's debt is SGP-noncompliant, the risk of sovereign default is higher, which creates an incentive for the Commission to actually enforce the SGP deficit criterion. The more severe the member state's noncompliance with the SGP debt criteria, the larger the decrease in an individual's support for monitoring. This is true regardless of an individual's support for the EMU.

Hypothesis 3. As a member state's noncompliance with the SGP deficit criterion increases, an individual's support for Commission monitoring will decrease, but only when the member state's debt is SGP-noncompliant.

3.4 Empirics

Studies of mass IPE depend on observational survey data, usually collected by someone other than the researcher (Urbatsch 2013; Curtis, Jupille and Leblang 2014; and Bearce and Tuxhorn 2017 use field original surveys; Hays, Ehrlich and Peinhardt 2005; Hainmueller and Hiscox 2006, 2007; Mansfield and Mutz 2009; and Pandya 2010 use published survey data). In this paper, I use Eurobarometer survey data. The Commission administers the Standard Eurobarometer twice per year, once in the spring and once in the fall.[7] Since the start of the Eurozone crisis, the Standard Eurobarometer has included a special bank of questions about attitudes towards specific economic policies. Rarely does a cross-national survey contain so many targeted questions about economic policy.

Each Standard Eurobarometer survey wave conducts face-to-face interviews with citizens of every EU member state. My sample includes data from waves 79 and 80, administered in the spring and fall of 2013, respectively. The bank of questions on the crisis changes from wave to wave, and these two waves are the only two that include questions on monitoring. Wave 79 includes the then-27 EU member states. Wave 80 adds Croatia, which joined the EU between the waves. I subset the sample to only include current Eurozone members plus Latvia and Lithuania.[8] In most member states, around a thousand respondents are interviewed, but in the smallest ones, that number is around five hundred. Since these two survey waves were administered only months apart, I pool them to construct a cross-sectional sample of 53,434 individual respondents.

I treat non-responses as missing values. The most common approach to missing data is list-wise deletion. Missing data due to item non-response, however, may impede inferences about the population if those non-responses are correlated with support for monitoring.

[7] I use the anonymized raw data provided by the GESIS – Leibniz Institute for the Social Sciences.

[8] I include Latvia and Lithuania because they were in the process of implementing the convergence criteria to join the EMU. Latvia joined on January 1, 2014 and Lithuania joined on January 1, 2015. Knowing they were set to join the EMU, respondents in these countries faced the same set of tradeoffs and incentives as those in current EMU members.

In many cases, the respondent only fails to answer one question needed to calculate the variables included in my models. I use multiple imputation to avoid throwing away data. I include a set of auxiliary variables to improve imputation of the dependent variable. I impute multiple data sets, perform analysis on each, combine the estimates, and apply the appropriate adjustments to the standard errors (Rubin 2004).

3.4.1 Measurement

The individual-level measures all come from the Eurobaromter data. These include measures of an individual's support for monitoring (*ex ante* and *ex post*), a measure of an individual's support for the EMU, and a set of individual-level control variables.

I use the following question to create individual-level measures of support for *ex ante* monitoring (i.e., pre-approval of member state budgets by the Commission) and *ex post* monitoring (i.e., financial sanctions for noncompliance): "A range of measures to tackle the current financial and economic crisis is being discussed in the European institutions. For each, could you tell me whether you think it would be effective or not?" The interviewer then shows the respondent a list of policies.

For each policy, the respondent can answer (1) "Very effective," (2) "Fairly effective," (3) "Not very effective," (4) "Not at all effective," or (5) "Don't know." One policy, corresponding to *ex ante* monitoring, is: "EU approval in advance of EU Member States' governments budgets." Another, corresponding to *ex post* monitoring, is: "Fines for EU Member States' governments that spend or borrow too much." I create a four-point index for each policy with higher values indicating stronger support.

I use the following question to create an individual-level measure of support for the EMU: "What is your opinion on each of the following statements? Please tell me for each statement, whether you are for it or against it." One statement is: "A European economic and monetary union with one single currency, the euro." Respondents can answer (1) "For," (2) "Against," or (3) “Don’t know.” I create a dummy variable that takes a value of one if the respondent supports the EMU.

To measure SGP-noncompliance, I use macroeconomic data from Eurostat. I measure noncompliance with the SGP debt criterion as the degree to which a member state’s debt exceeds the SGP criterion of 60 percent of GDP, expressed as a percent of GDP. This measure ranges from 0 percent of GDP for all member states that are in compliance to 99.6 percent of GDP. Only 47.51 percent of individuals live in a member state that are in compliance with the SGP debt criterion. Unsurprisingly, Greece’s is the maximum value: its sovereign debt exceeds the SGP limit by an enormous 99.6 percent of GDP.

I measure noncompliance with the SGP deficit criterion using a dummy variable that indicates whether a member state’s deficit exceeds the SGP criterion of 3 percent of GDP, and therefore whether the Commission has the option to impose austerity measures on the member state. Only 37.87 percent of individuals in the sample live in a member state that are in compliance with the SGP deficit criterion. For both measures, I use data from 2012 to avoid post-treatment bias.

As previously mentioned, I impute missing values. The two measures of noncompliance with the SGP criteria have no missingness. Support for the EMU is missing 6.59 percent of observations (3,302). Support for *ex ante* monitoring is missing 14.24 percent of observations (6,659), whereas support for *ex post* monitoring is missing 8.23 percent (4,057).[9]

I control for a variety of individual-level factors, including: whether an individual views the EU institutions as competent to respond to the crisis (a dummy); an individual’s sophistication (based on her factual knowledge of the EU); an individual’s cosmopolitanism (based on how frequently she exhibits a set of cosmopolitan behaviors); an individual’s household financial situation; an individual’s level of education; an individual’s gender; and an indi-

[9] The imputation model includes the two dependent variables, the three independent variables of interest, the control variables, and a set of auxiliary variables used in the matching model described below.

vidual's age. Data on these respondent-level measures come from the same Eurobarometer surveys described above.

I also control for two country-level factors: the severity of a member state's austerity policies and its unemployment rate. Austerity refers to measures that reduce a member state's deficit. I operationalize austerity as the percent change in a member state's deficit from the onset of the sovereign debt crisis in 2009 to the administration of the survey in 2013. Data for these country-level measures come from Eurostat.

3.4.2 Descriptive Statistics

The respondent-level measures vary considerably across member states. Figure 3.1 displays descriptive statistics on the key respondent-level measures: support for *ex ante* monitoring, support for *ex post* monitoring, and support for the EMU. Panels A and B describe support for *ex ante* monitoring. In all eurozone members, a majority of respondents supports pre-approval of member state budgets. In each member state, between 55 percent and 80 percent of respondents answer that they think pre-approval of member states budgets is "Fairly effective" or "Very effective." Pooling respondents across member states, the modal response was "Fairly effective." Panels C and D describe support for *ex post* monitoring. Again, support is high: between 55 percent and 90 percent of respondents support monitoring. The modal category is again "Fairly effective," but there are far more responses of "Very effective." Panels E and F describe support for the EMU. Across member states, between 45 percent and 85 percent of respondents support the EMU. Overall, respondents support the EMU by a two-to-one margin.

These statistics indicate that respondents think *ex post* monitoring provides a better solution to the stability/austerity tradeoff. It could be that less people support *ex ante* monitoring because it is a tighter control mechanism than *ex post* monitoring, making it easier for the Commission to incentivize costly austerity measures.

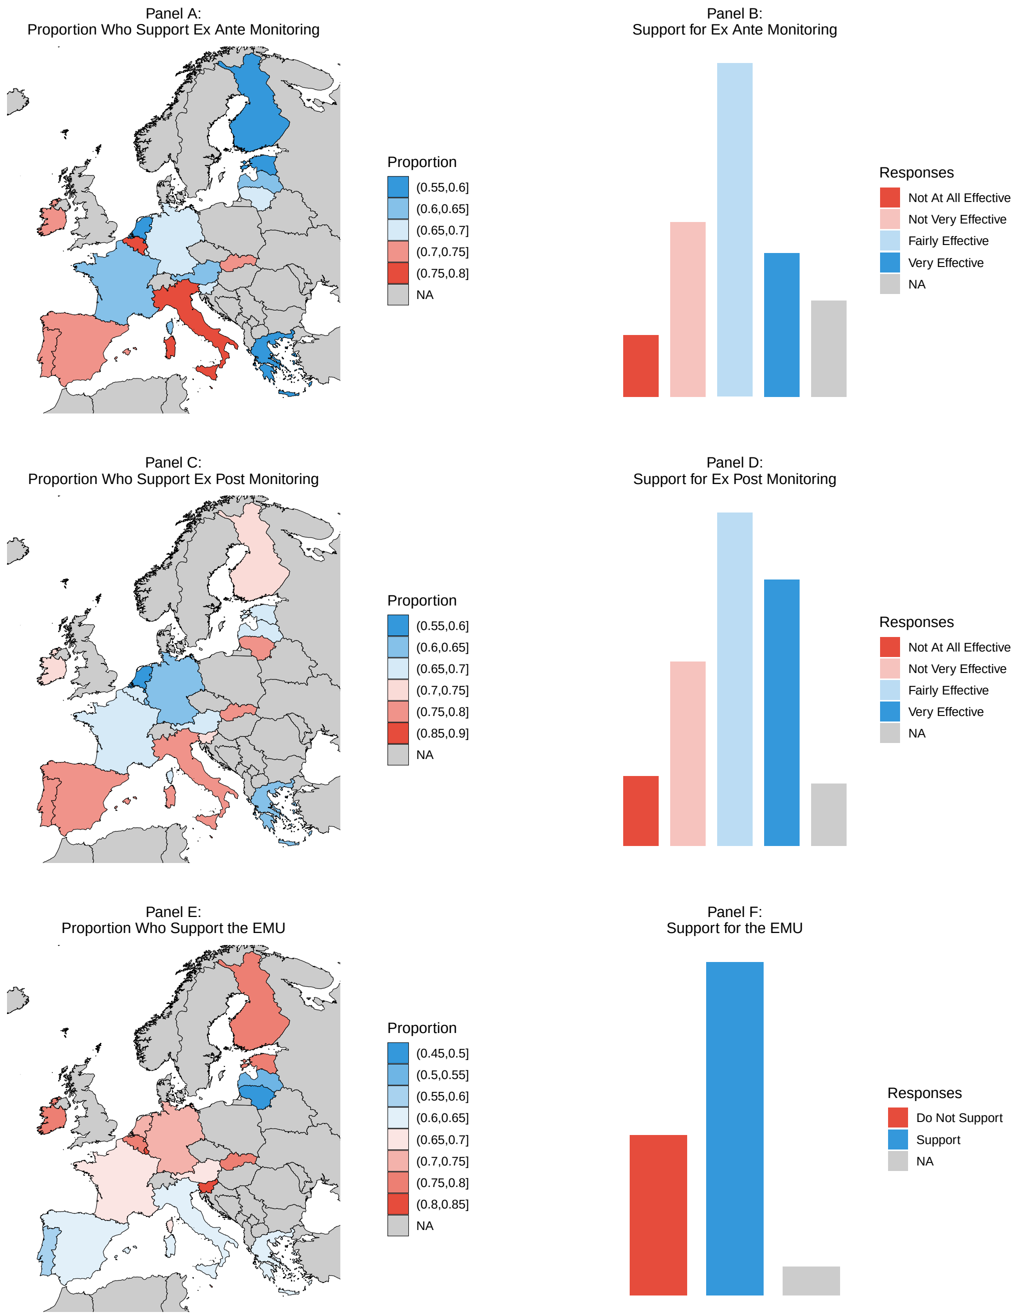

Figure 3.1. Descriptive Statistics for Respondent-Level Measures

Note: This figure shows descriptive statistics for the respondent-level measures. Panel A shows the proportion of respondents who think *ex ante* monitoring will be fairly effective or very effective by Eurozone member. Panel B shows the distribution of responses with respect to *ex ante* monitoring across all Eurozone members. Panel C shows the proportion of respondents who think *ex post* monitoring will be fairly effective or very effective by Eurozone member. Panel D shows the distribution of responses with respect to *ex post* monitoring across all Eurozone members. Panel E shows the proportion of respondents who support the EMU by Eurozone member. Panel F shows the proportion of respondents across all Eurozone members that support the EMU.

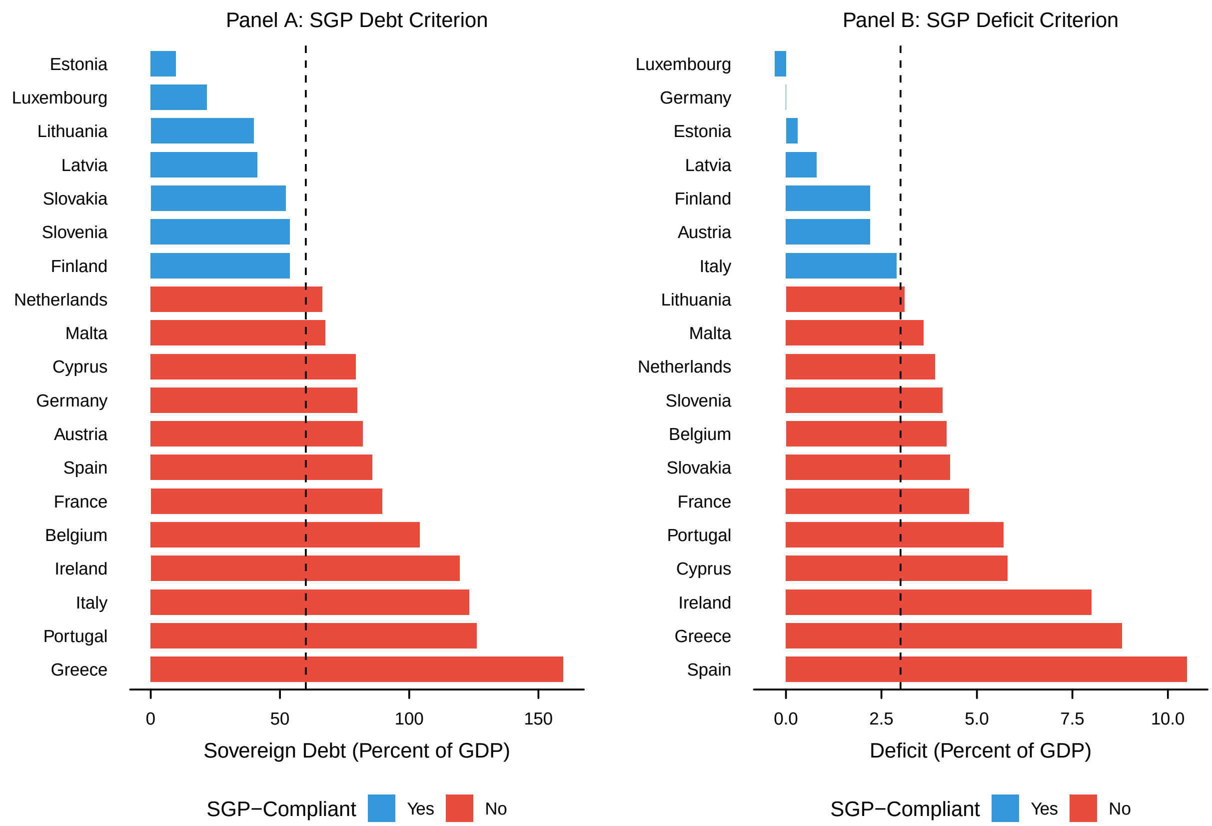

Figure 3.2. Noncompliance with the SGP Criteria

Note: This figure shows noncompliance with the SGP debt criterion (debt as a percent of GDP in excess of 60 percent) and with the SGP deficit criterion (deficit as a percent of GDP in excess of 3 percent) in 2012. Germany did not run a deficit or a surplus in 2012.

Figure 3.2 shows descriptive statistics on the key country-level measures: noncompliance with the SGP's debt and deficit criteria (at the end of 2012). Panel A shows noncompliance with the debt criterion (60 percent of GDP), whereas Panel B shows noncompliance with the deficit criterion (3 percent of GDP). What is striking is the extensiveness and severity of noncompliance. Only 7 of the 19 member states in the sample are in compliance with the debt criterion, and that includes Latvia and Lithuania, whose upcoming membership was conditional on full compliance with the SGP criteria.

Compliance with the deficit criterion is no better. Again, only 7 of the 19 member states in the sample are in compliance (although the 7 differ). Critically, both France (debt and deficit) and Germany (debt) are in violation of the SGP criteria. Their noncompliance disincentivized the Commission from enforcing the SGP, opening the door for others to follow (Baerg and Hallerberg 2016).

3.4.3 Estimation Strategy

The structure of the data is multilevel: individuals are nested within states. To account for this structure, I estimate multilevel models (Gelman 2006; Armingeon and Ceka 2014; Gomez 2015). Multilevel modeling offers two main advantages. First, it accounts for individual-level and member state-level variation in estimating the member state-level regression coefficients. Second, it allows me to take into account member state-level variation in the uncertainty over individual-level coefficients (Gelman 2006).

I estimate varying-intercept multilevel linear regression models that include individual-level predictors, member state-level predictors, and cross-level interaction terms. This estimator allows the intercept to vary by member state but not the slopes of individual-level predictors. A varying-intercept, varying-slope model would estimate the variance of the slope across member states (Gelman 2006). There is no *a priori* theoretical basis for assuming that the slopes of individual-level predictors of interest vary across member states, apart from the interaction effects that I hypothesize and model.

The sample includes $i = 1, \ldots, n$ individuals nested in $j = 1, \ldots, J$ member states. Let $j[i]$ be the member state in which individual i is located. There are $l = 1, \ldots, L$ policies that an individual i has preferences over. The model includes individual-level predictors x_i and member state-level predictors x_j. Including cross-level interaction terms, there are k predictors. Thus, including a constant, there are $k + 1$ parameters to estimate. I estimate the following varying-intercept model:

$$y_i^l = N(\alpha_{j[i]} + X_i\beta, \sigma^2_{y^l}), \quad \text{for } i = 1, \ldots, n \tag{3.1}$$

$$\alpha_j \sim N(\mu_\alpha, \sigma^2_\alpha), \tag{3.2}$$

where $\alpha_{j[i]}$ is a constant and X is a $n \times k$ matrix composed of column vectors for individual-level predictors, member state-level predictors, and in some specifications, same-level or cross-level interaction terms, but not a constant. Equation (3.1) is the lower level and Equation (3.2) is the upper level. Since this is a varying-intercept model, the only coefficient that varies across member states is the constant, $\alpha_{j[i]}$. We can equivalently write the model with normally distributed member state-specific errors, η_j, instead of with member state-specific intercepts with normally distributed errors:

$$y_i^l = N(X_i\beta + \eta_{j[i]}, \sigma^2_{y^l}), \quad \text{for } i = 1, \ldots, n \tag{3.3}$$

$$\eta_j \sim N(0, \sigma^2_\alpha), \tag{3.4}$$

where X is an $n \times (k+1)$ matrix composed of column vectors for a constant, individual-level predictors, member state-level predictors, and cross-level interaction terms. Again, Equation (3.3) is the lower level and Equation (3.4) is the higher level.

I estimate multilevel linear models instead of multilevel ordered logit models. Honaker, King and Blackwell (2011) recommend imputing ordinal variables as if they were continuous, as non-integer values contain information that is if we force-imputed values to be integers. Moreover, it is easier to interpret marginal effects with linear models when there are interaction terms, which requires marginal effects plots (Brambor, Clark and Golder 2006; Berry, Golder and Milton 2012). With ordered logit models, we have to estimate the marginal effect on the probability of observing each ordered category.

3.4.4 Analysis and Findings

This section presents the results from the imputed sample. The non-imputed results are substantively similar. I find empirical support for each hypothesis. Across the board, the sizes of the substantive effects of the variables of interest are small. This is not surprising, however, because we should expect public opinion data on such a complex topic to be extremely noisy. As many scholars have found (e.g., Curtis, Jupille and Leblang 2014), non-opinions dominate. And even during a financial crisis, salience is only so high. The substantive effects of the control variables in the model are also small.

I estimate a three-way interaction between support for the EMU, noncompliance with the SGP deficit criterion, and noncompliance with the SGP debt criterion. The three-way interaction term has a statistically significant effect on both types of monitoring. Likelihood-ratio tests indicate these multilevel models provide a better fit than OLS models. To evaluate the hypotheses, we need to examine marginal effects plots. Figure 3.3 shows 6 marginal effects plots. Panels A and B test Hypotheses 1a and 1b; Panels C and D test Hypotheses 2a and 2b; and Panels E and F test Hypothesis 3.

Hypothesis 1a predicts that the marginal effect of noncompliance with the debt criterion will be positive when an individual supports the EMU and her member state is complying with the deficit criterion. Hypothesis 1b, on the other hand, predicts that the marginal effect will be negative when and individual does not support the EMU and her member state is not complying with the deficit criterion. Theory does not make a prediction for the other two cases. The results in Panel A, with respect to *ex ante* monitoring, are consistent with these predictions. A marginal increase in noncompliance with the debt criterion has a statistically significant positive effect (0.006, $p < 0.01$) when the respondent supports the EMU and the respondent's member state is complying with the deficit criterion and a statistically significant negative effect (-0.003, $p < 0.01$) when the respondent does not support the EMU and the respondent's member state is not complying with the deficit criterion. The results in Panel B, with respect to *ex post* monitoring, are substantively similar (0.003, $p < 0.05$; -0.005, $p < 0.01$).

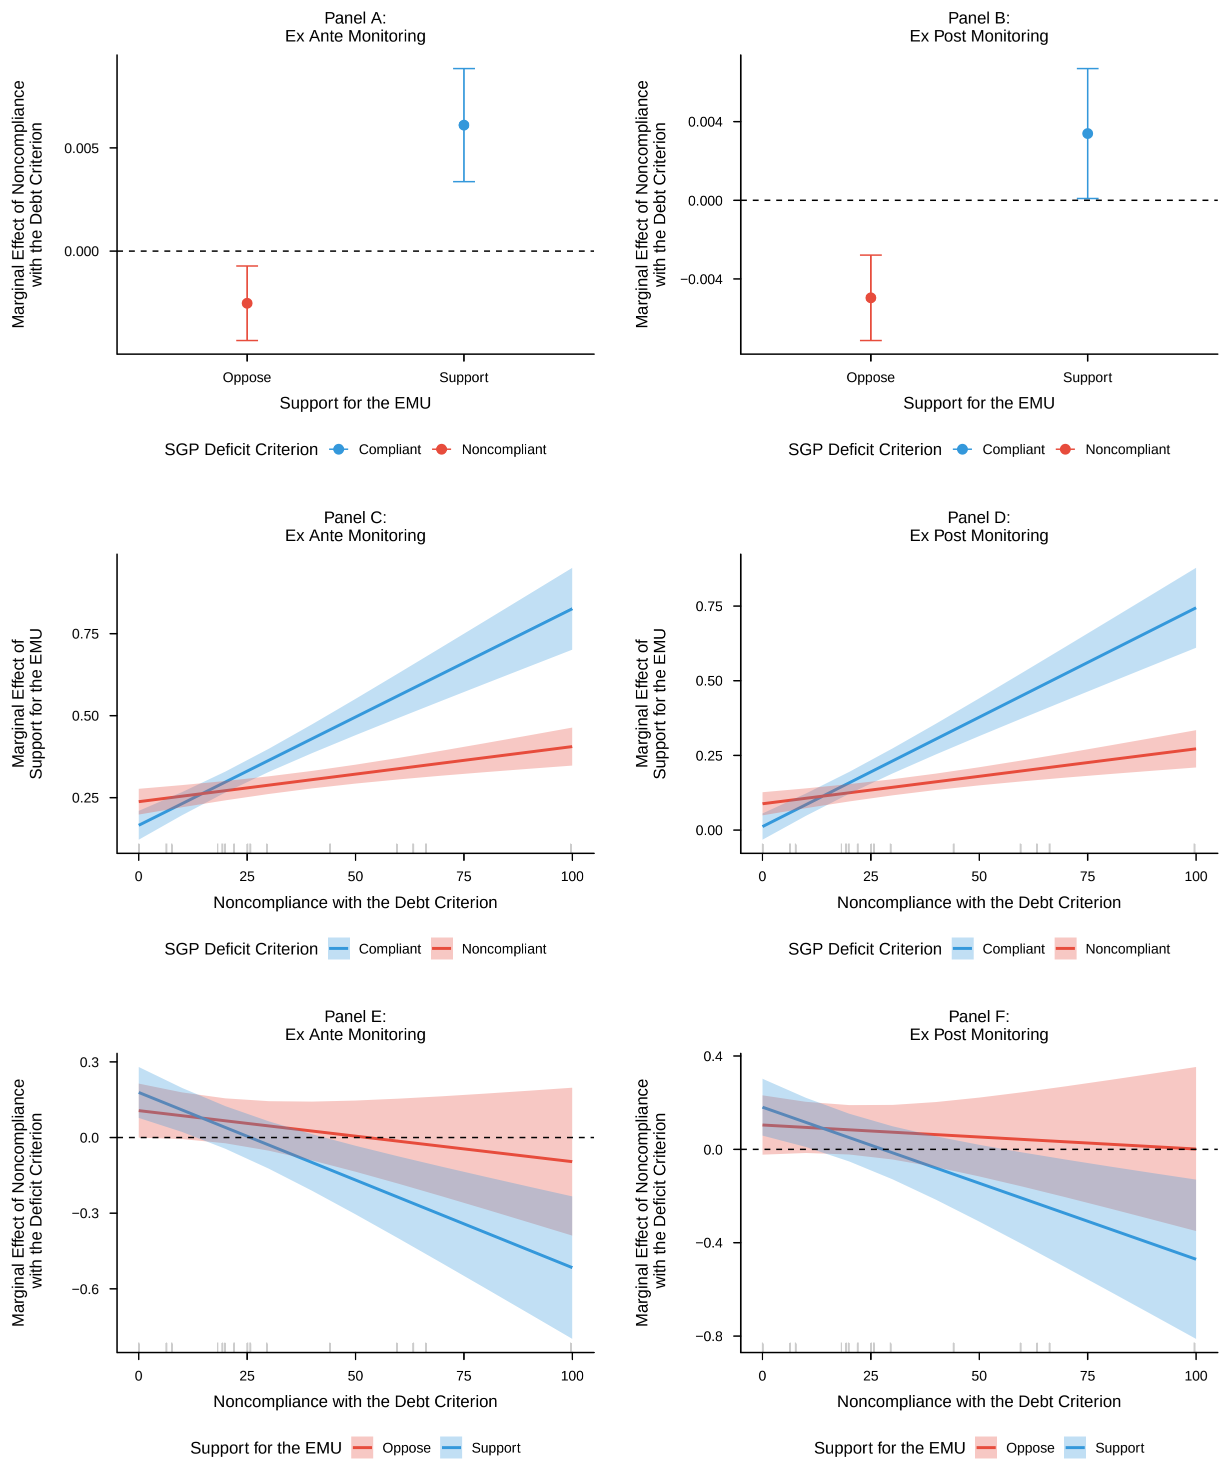

Figure 3.3. Marginal Effect Plots (Imputed Sample)

Note: This figure shows the results from the imputed sample. It shows the conditional marginal effects of noncompliance with the SGP debt criterion (Panels A and B), support for the EMU (Panels C and D), and noncompliance with the SGP deficit criterion (Panels E and F) on support for *ex ante* monitoring (Panels A, C, and E) and *ex post* monitoring (Panels B, D, and F).

Hypothesis 2a predicts that the marginal effect of support for the EMU will be positive, regardless of a member state's noncompliance with the SGP criteria. Hypothesis 2b predicts that the effect will be smallest with a member state's debt is SGP-complaint and largest when its deficit is SGP-compliant and its debt is SGP-noncompliant. Panels C and D provide clear support for this prediction. Consistent with Hypothesis 2a, the marginal effect is always positive. The effect is smallest when the member state's debt is SGP-compliant. The size of the effect is substantively similar regardless of the member state's compliance with the SGP deficit criterion. The effect is largest when the member state's deficit is SGP-compliant and its noncompliance with the SGP debt criterion is severe. This is consistent with Hypothesis 2b. In Panel C, the estimated effect increases from 0.16 to 0.83 as noncompliance with the debt criterion increases (over its in-sample range), conditional on the member state's deficit being SGP-compliant. When the member state's deficit is SGP-noncompliant, the effect ranges from 0.23 to 0.41. The effects in Panel D are substantively similar.

Hypothesis 3 predicts that the marginal effect of a member state's noncompliance with the deficit criterion will be negative, but only when the member state's debt is SGP-noncompliant. Panels C and D provide mixed support for this hypothesis. The marginal effect of an increase in the member state's noncompliance with the SGP deficit criterion is negative when the member state's debt is sufficiently SGP-noncompliant, but only if the respondent supports the EMU. For respondents who support the EMU, the negative effect is statistically significant for a sufficiently high level of noncompliance with the SGP debt criterion (over 50 percent of GDP). However, the effect is not statistically significant for respondents who oppose the EMU, although the sign is negative for *ex ante* monitoring (see Panel E). As predicted, the effect is insignificant when the respondent's member state's noncompliance with the debt criterion is low, as the Commission is less likely to enforce relatively minor violations.

Looking at Panel F, noncompliance with the SGP deficit criterion only has a statistically significant effect on support for *ex post* monitoring for the largest values of noncompliance with the SGP debt criterion, suggesting that respondents only believe that the Commission will enforce the SGP deficit criterion, thereby imposing austerity, when a member state's noncompliance with the debt criterion is extremely poor (i.e., the member state is running a real risk of sovereign default). This is not surprising. The Commission's self-selection out of *ex post* monitoring directly contributed to the outbreak of the sovereign debt crisis. As such, a belief that the Commission would only enforce the SGP deficit criterion for the most at-risk member states would be well-founded.

Overall, I find evidence that, in the aggregate, the preferences of survey respondents are largely consistent with individual rationality. Consistent with theory, support for the EMU interacts with a member state's compliance with the SGP criteria to shape support for *ex ante* and *ex post* Commission monitoring. The evidence does not indicate whether individuals rationally evaluate the policy of Commission monitoring, however.

3.4.5 Challenges to Inference

An important inferential challenge is that a respondent's support for the EMU is not randomly assigned. I black-box the origins of an individual's self-reported support for the EMU and implicitly treat that variable as exogenous, but my empirical analysis needs to address the possibility that an individual's self-reported support for the EMU is endogenous to the public discourse about possible policy responses to the Eurozone crisis.

For my estimates to be unbiased, support for the EMU needs to be distributed as if at random, conditional on the covariates in the model, but there could be systematic differences between people who support the EMU and those that do not along dimensions that predict support for Commission monitoring. To increase confidence that this form of endogeneity is not significantly biasing my estimates, I use matching as a preprocessing technique and re-run my analysis using the matched sample (Ho et al. 2007). This estimates the average treatment effect on the treated (ATT), which is the average effect of support for the EMU on support for monitoring among respondents who support the EMU.

Note that this matching design does not address another source of endogeneity. One of the independent variables of interest (whether an individual supports the EMU) is endogenous to the dependent variable (support for Commission monitoring of member state compliance with the SGP criteria) because supporting the EMU implicitly means supporting the institutional structure of the monetary union, which includes the SGP and Commission monitoring. In other words, Commission monitoring of the SGP criteria could influence an individual's support for the EMU. However, among drivers of public support for the EMU, the institutional structure of the monetary union is unlikely to rank highly.

I use one-to-one propensity score matching (Rosenbaum and Rubin 1983). Each treated observation is matched to a control observation based on its propensity of getting the treatment. This process approximates an experiment in the sense that each member in a pair of matched observations have the same likelihood of getting the treatment, but only one actually does. This eliminates extreme counterfactuals (King and Zeng 2006). I match each treated observation to the control observation with the closest propensity score (i.e., nearest neighbor matching). Matching requires common support; that is, the distributions of propensity scores for the treatment and control groups need to overlap sufficiently. To increase the number of treated observations in the region of common support, I match with replacement using a caliper.

There are two common strategies to implementing propensity score techniques using multilevel data: within-cluster matching and across-cluster matching (Kim and Steiner 2015). The within-cluster approach estimates propensity scores by stratum (here, by member state) and then only matches treatment and control observations that come from the same stratum. As such, this approach only uses individual-level covariates in the matching model (i.e., the model used to generate the propensity scores). The biggest disadvantage with the within-cluster approach is that it can be harder to find good matches when matched pairs must come from the same stratum. The across-cluster approach incorporates stratum-level measures into the matching model and matches across strata. The strata, then, only matter

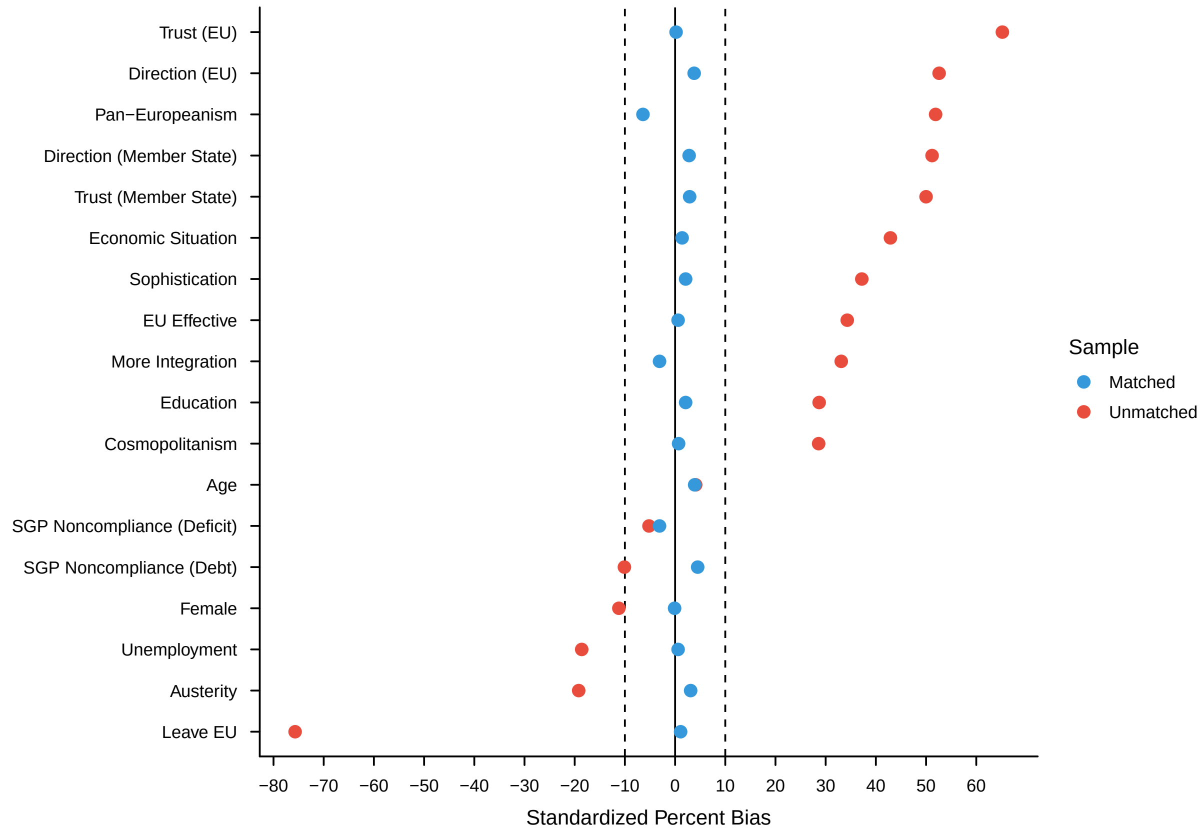

Figure 3.4. Balance Plot

Note: This figure shows the balance between the matched and unmatched samples for all covariates in the matching model.

to the matching algorithm insofar as they affect an observation's propensity score. To get the closest matches, I use across-cluster matching.[10]

Matching on observables is appropriate when the factors that determine treatment assignment are observable and measurable. The matching model (i.e., the model used to generate the propensity scores) includes my control variables and several auxiliary variables that I expect to be related to the selection mechanism.

[10] To increase the number of treated observations in the region of common support, I match with replacement using a caliper. This allows treatment observations to be paired with multiple control observations or *vice versa*, but only if their propensity scores are sufficiently close. The matching algorithm assigns each observation a weight that indicates how many times it is used to construct a matched pair. The vast majority of observations are only used once.

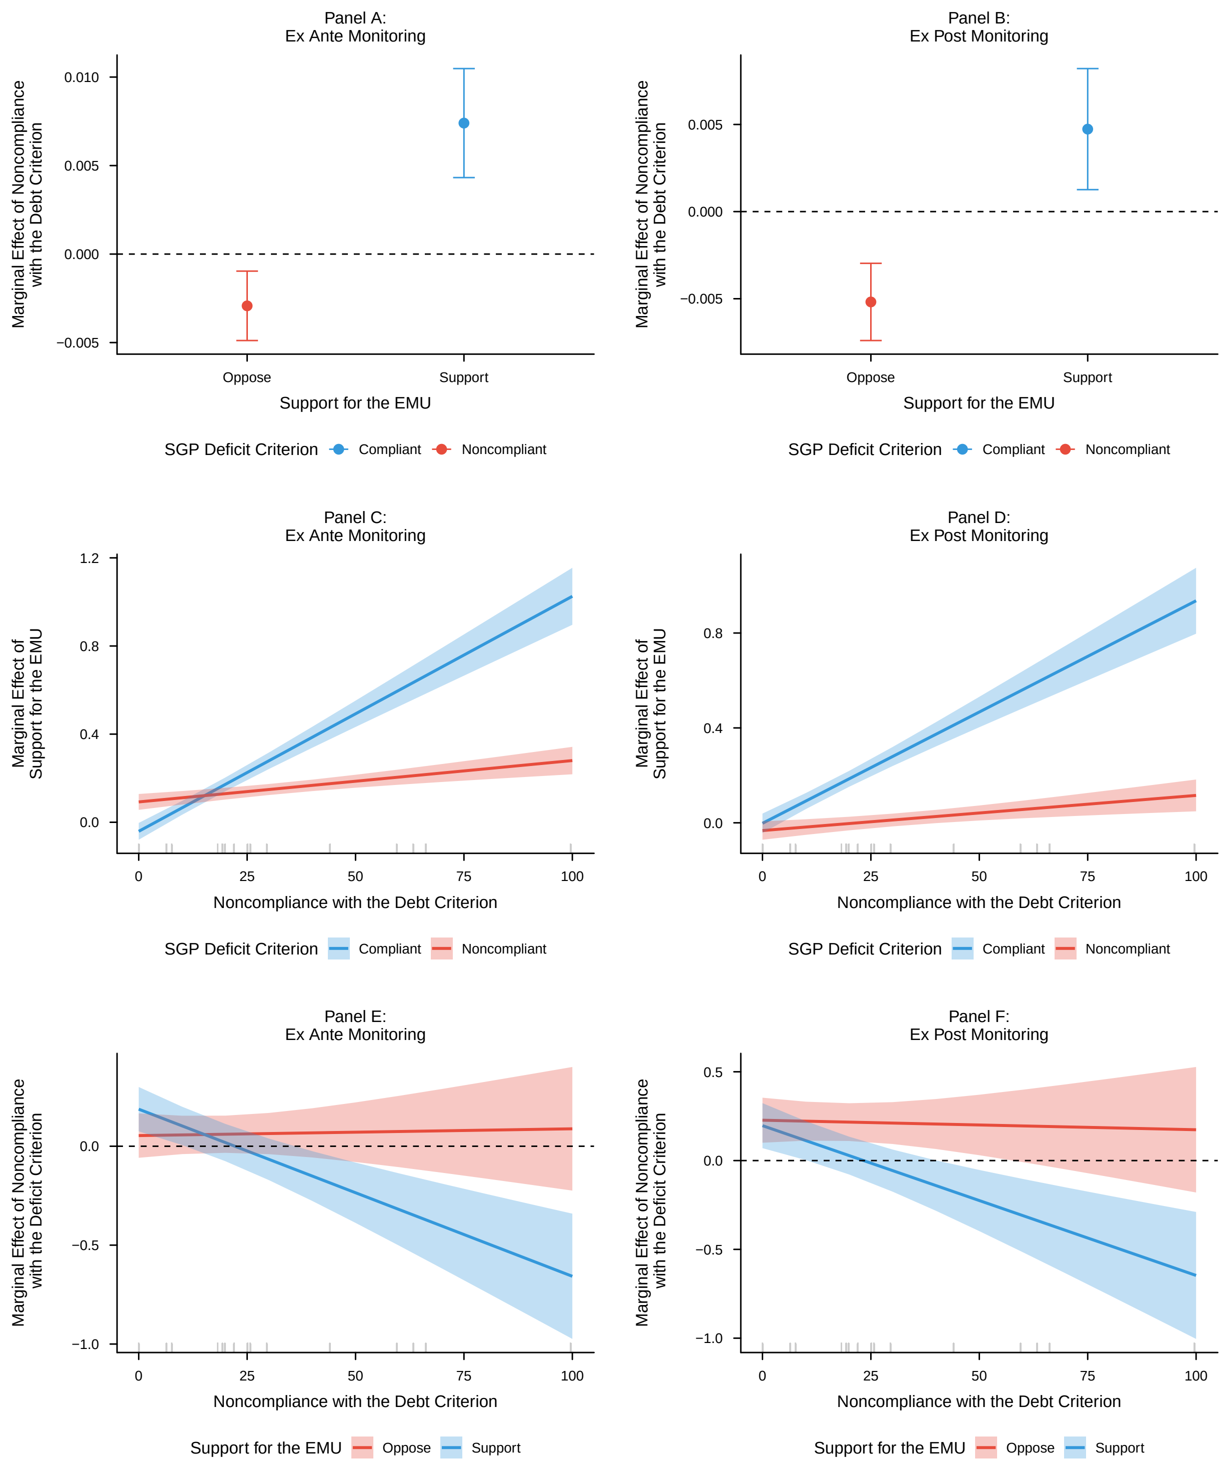

Figure 3.5. Marginal Effect Plots (Matched Sample)

Note: This figure shows the results from the matched sample. It shows the conditional marginal effects of noncompliance with the SGP debt criterion (Panels A and B), support for the EMU (Panels C and D), and noncompliance with the SGP deficit criterion (Panels E and F) on support for *ex ante* monitoring (Panels A, C, and E) and *ex post* monitoring (Panels B, D, and F).

The first set of measures capture ideology. Many studies have identified a pro-/anti-EU dimension to political ideology in Europe. Recent research has shown that framing effects are critical in the design of survey questions, so I include two measures of pro-/anti-EU ideology (Hiscox 2006; Ardanaz, Murillo and Pinto 2013). One is based on a question framed around support for further EU integration, the other based on a question framed around support for leaving the EU. A second set of measures capture trust in EU institutions and member state institutions, as an individual's overall attitudes towards these institutions could shape which she thinks should implement monetary policy. A third set of measures capture an individual's general assessment of the direction in which the EU and her member state are moving. A final measure captures an individual's pan-Europeanism (i.e., the degree to which she identifies as European versus her nationality), which could be correlated with affinity for the euro as a symbol of European cooperation.

To assess the balance of the treatment and control groups, I calculate the standardized percent bias. A common rule of thumb is that the standardized percent bias should be less than 25 percent for all covariates in the matching model (Ho et al. 2007). As shown in Figure 3.4, the matched sample is far more balanced than the unmatched sample: the standardized percent bias of all covariates is less than 10 percent.

Finally, I re-run my analysis using the matched sample. All of the results are substantively the same as those from my analysis using the unmatched sample (see Figure 3.5). These results indicate that, conditional on the matching model including the appropriate confounding variables, that respondents' self-reported support for the EMU is not endogenous to their support for monitoring.

3.5 Conclusion

This paper goes in a different direction than recent literature by looking for evidence of rationality (at the aggregate level) in the public's evaluation of complex economic policies.

Recent studies think about rationality in terms of whether well-informed individuals are more likely to support policies that advance their material interests (e.g., Gomez and Wilson 2006; Bearce and Tuxhorn 2017), but this paper thinks about rationality in terms of whether policy preferences, in the aggregate, take into account context-specific tradeoffs that international economic policies present, conditional on national economic conditions, and through the lens of their own self-defined interests. In light of a recent global trend toward using direct democracy to make major economic policy decisions, this question has important implications for the coherence of economic policy.

I provide evidence from the European sovereign debt crisis that, in a high-information environment, aggregate preferences are surprisingly consistent with individual rationality. Specifically, I look at Commission monitoring of member state compliance with the SGP. I find that evidence that individuals (a) can identify the likely outcomes that a policy can generate, (b) understand how the current economic situation in their country conditions the effects of a policy, (c) evaluate any tradeoffs that a policy creates, and (d) support policies that advance their self-defined interests, on balance.

Scholars can use these criteria to assess how respondents evaluate complex, multi-dimensional policies in other contexts. Beyond establishing the generality of these findings, future research should look at how the information environment impacts respondent sophistication. I choose the European sovereign debt crisis as the substantive context for this paper specifically because it is a high-information environment. In a high-information environment where the policy in question is the subject of intense, sustained media coverage, respondents are more likely to be familiar with the economic tradeoffs. An important next step is to learn whether these findings generalize to contexts in which rationality requires more respondents to independently seek out information and do more of the cognitive work themselves without elite cues from politicians and the media.

Bibliography

Alter, Karen. 2001. *Establishing the Supremacy of European Law: The Making of an International Rule of Law in Europe.* Oxford, England: Oxford University Press.

Alter, Karen J. 2000. "The European Union's Legal System and Domestic Policy: Spillover or Backlash?" *International Organization* 54(3):489–518.

Ardanaz, Martin, M. Victoria Murillo and Pablo M. Pinto. 2013. "Sensitivity to Issue Framing on Trade Policy Preferences: Evidence from a Survey Experiment." *International Organization* 67(02):411–437.

Armingeon, Klaus and Besir Ceka. 2014. "The Loss of Trust in the European Union During the Great Recession Since 2007: The Role of Heuristics from the National Political System." *European Union Politics* 15(1):82–107.

Axtell, Robert L. 2001. "Zipf Distribution of US Firm Sizes." *Science* 293(5536):1818–1820.

Baerg, Nicole Rae and Mark Hallerberg. 2016. "Explaining Instability in the Stability and Growth Pact: The Contribution of Member State Power and Euroskepticism to the Euro Crisis." *Comparative Political Studies* 49(7):968–1009.

Bearce, David H. and Kim-Lee Tuxhorn. 2017. "When are Monetary Policy Preferences Egocentric? Evidence from American Surveys and an Experiment." *American Journal of Political Science* 61(1):178–193.

Berry, William D., Matt Golder and Daniel Milton. 2012. "Improving Tests of Theories Positing Interaction." *The Journal of Politics* 74(3):653–671.

Blauberger, Michael and Susanne K. Schmidt. 2017. "The European Court of Justice and Its Political Impact." *West European Politics* 40(4):907–918.

Börzel, Tanja A. 2001. "Non-Compliance in the European Union: Pathology or Statistical Artefact?" *Journal of European Public Policy* 8(5):803–824.

Börzel, Tanja A. 2003. Guarding the Treaty: The Compliance Strategies of the European Commission. In *The State of the European Union*, ed. Tanja A. Börzal and Rachel A. Cichowski. Oxford: Oxford University Press pp. 197–220.

Brambor, T., WIlliam Roberts Clark and Matt Golder. 2006. "Understanding Interaction Models: Improving Empirical Analyses." *Political Analysis* 14(1):63–82.

Burley, Anne-Marie and Walter Mattli. 1993. "Europe Before the Court: A Politial Theory of Legal Integration." *International Organization* 47(Winter):41–76.

Carpenter, Bob, Andrew Gelman, Matthew D. Hoffman, Daniel Lee, Ben Goodrich, Michael Betancourt, Marcus A. Brubaker, Jiqiang Guo, Peter Li and Allen Riddell. 2017. "Stan: A Probabilistic Programming Language." *Journal of Statistical Software* 76(1).

Carrubba, Cliff, Barry Friedman, Andrew D. Martin and Georg Vanberg. 2012. "Who Controls the Content of Supreme Court Opinions?" *American Journal of Political Science* 56(2):400–412.

Carrubba, Clifford J. 2005. "Courts and Compliance in International Regulatory Regimes." *The Journal of Politics* 67(3):669–689.

Carrubba, Clifford J. 2009. "A Model of the Endogenous Development of Judicial Institutions in Federal and International Systems." *The Journal of Politics* 71(1):55–69.

Carrubba, Clifford J. and Anand Singh. 2004. "A Decision Theoretic Model of Public Opinion: Guns, Butter, and European Common Defense." *American Journal of Political Science* 48(2):218–231.

Carrubba, Clifford J. and Christopher Zorn. 2010. "Executive Discretion, Judicial Decision Making, and Separation of Powers in the United States." *The Journal of Politics* 72(3):812–824.

Carrubba, Clifford J. and Matthew Gabel. 2017. "International Courts: A Theoretical Assessment." *Annual Review of Political Science* 20:55–73.

Carrubba, Clifford J., Matthew Gabel and Charles Hankla. 2008. "Judicial Behavior under Political Constraints: Evidence from the European Court of Justice." *American Political Science Review* 102(4):435–452.

Carrubba, Clifford J. and Matthew J. Gabel. 2015. *International Courts and the Performance of International Agreements*. New York: Cambridge University Press.

Castro-Montero, José Luis, Edwin Alblas, Arthur Dyevre and Nicolas Lampach. 2018. "The Court of Justice and Treaty Revision: A Case of Strategic Leniency?" *European Union Politics* 19(4):570–596.

Chaney, Thomas. 2008. "Distorted Gravity: The Intensive and Extensive Margins of International Trade." *American Economic Review* 98(4):1707–1721.

Chayes, Abram and Antonia Handler Chayes. 1993. "On Compliance." *International Organization* 47(2):175–205.

Conant, Lisa J. 2002. *Justice Contained: Law and Politics in the European Union*. Cornell University Press.

Curtis, K. Amber, Joseph Jupille and David Leblang. 2014. "Iceland on the Rocks: The Mass Political Economy of Sovereign Debt Resettlement." *International Organization* 68(3):721–740.

Demidova, Svetlana and Andrés Rodríguez-Clare. 2009. "Trade Policy Under Firm-Level Heterogeneity in a Small Economy." *Journal of International Economics* 78(1):100–112.

Dixit, Avinash K. and Joseph E. Stiglitz. 1977. "Monopolistic Competition and Optimum Product Diversity." *American Economic Review* 76(3):297–308.

Downs, George W., David M. Rocke and Peter N. Barsoom. 1996. "Is the Good News about Compliance Good News about Cooperation?" *International Organization* 50(3):379–406.

Fjelstul, Joshua C. and Clifford J. Carrubba. 2018. "The Politics of International Oversight: Strategic Monitoring and Legal Compliance in the European Union." *American Political Science Review* 112(3):429–445.

Fordham, Benjamin O. and Katja B. Kleinberg. 2012. "How Can Economic Interests Influence Support for Free Trade?" *International Organization* 66(2):311–328.

Franzese, Robert J. and Cindy Kam. 2009. *Modeling and Interpreting Interactive Hypotheses in Regression Analysis.* University of Michigan Press.

Gabaix, Xavier. 2009. "Power Laws in Economics and Finance." *Annual Review of Economics* 1(1):255–294.

Gabel, Matthew. 2000. "European Integration, Voters and National Politics." *West European Politics* 23(4):52–72.

Garrett, Geoffrey, R. Daniel Kelemen and Heiner Schulz. 1998. "The European Court of Justice, National Governments, and Legal Integration in the European Union." *International Organization* 52(1):149–176.

Gelman, Andrew. 2006. "Prior Distributions for Variance Parameters in Hierarchical Models." *Bayesian Analysis* 1(3):515–534.

Gelman, Andrew and Jennifer Hill. 2006. *Data Analysis Using Regression and Multilevel/Hierarchical Models*. Cambridge University Press.

Gilligan, Michael J. and Leslie Johns. 2012. "Formal Models of International Institutions." *Annual Review of Political Science* 15(7):1–23.

Gilligan, Michael, Leslie Johns and B. Peter Rosendorff. 2010. "Strengthening International Courts and the Early Settlement of Disputes." *Journal of Conflict Resolution* 54(1):5–38.

Gomez, Brad T. and J. Matthew Wilson. 2006. "Cognitive Heterogeneity and Economic Voting: A Comparative Analysis of Four Democratic Electorates." *American Journal of Political Science* 50(1):127–145.

Gomez, Raul. 2015. "The Economy Strikes Back: Support for the EU During the Great Recession." *Journal of Common Market Studies* 53(3):577–592.

Grossman, Gene and Elhanan Helpman. 1994. "Protection for Sale." *American Economic Review* 84(4):833–850.

Hainmueller, Jens and Michael J. Hiscox. 2006. "Learning to Love Globalization: Education and Individual Attitudes Toward International Trade." *International Organization* 60(2):469–498.

Hainmueller, Jens and Michael J. Hiscox. 2007. "Educated Preferences: Explaining Attitudes Toward Immigration in Europe." *International Organization* 61(2):399–442.

Hainmueller, Jens and Michael J. Hiscox. 2010. "Attitudes Toward Highly Skilled and Low-Skilled Immigration: Evidence from a Survey Experiment." *American Political Science Review* 104(1):61–84.

Hartlapp, Miriam and Gerda Falkner. 2009. "Problems of Operationalization and Data in EU Compliance Research." *European Union Politics* 10(2):281–304.

Hays, Jude C., Sean D. Ehrlich and Clint Peinhardt. 2005. "Government Spending and Public Support for Trade in the OECD: An Empirical Test of the Embedded Liberalism Thesis." *International Organization* 59(2):473–494.

Helpman, Elhanan, Marc J. Melitz and Stephen R. Yeaple. 2004. "Export Versus FDI with Heterogeneous Firms." *American Economic Review* 94(1):300–316.

Hicks, Raymond, Helen V. Milner and Dustin Tingley. 2014. "Trade Policy, Economic Interests, and Party Politics in a Developing Country: The Political Economy of CAFTA-DR." *International Studies Quarterly* 58(1):106–117.

Hiscox, Michael J. 2006. "Through a Glass and Darkly: Attitudes Toward International Trade and the Curious Effects of Issue Framing." *International Organization* 60(3):755–780.

Ho, Daniel E., Kosuke Imai, Gary King and Elizabeth A. Stuart. 2007. "Matching as Nonparametric Preprocessing for Reducing Model Dependence in Parametric Causal Inference." *Political Analysis* 15(3):199–236.

Hoffman, Matthew D. and Andrew Gelman. 2014. "The No-U-Turn Sampler: Adaptively Setting Path Lengths in Hamiltonian Monte Carlo." *Journal of Machine Learning Research* 15:1593–1623.

Honaker, James, Gary King and Matthew Blackwell. 2011. "Amelia II: A Program for Missing Data." *Journal of Statistical Software* 45(7):1–47.

Jensen, Christian. B. 2007. "Implementing Europe: A Question of Oversight." *European Union Politics* 8(4):451–477.

Johns, Leslie. 2012. "Courts as Coordinators: Endogenous Enforcement and Jurisdiction in International Adjudication." *Journal of Conflict Resolution* 56(2):257–289.

Johns, Leslie. 2015. *Strengthening International Courts: The Hidden Costs of Legalization*. Michigan Studies in International Political Economy University of Michigan Press.

Jupille, Joseph and David Leblang. 2007. "Voting for Change: Calculation, Community, and Euro Referendums." *International Organization* 61(4):763–782.

Keohane, Robert. 1984. *After Hegemony: Cooperation and Discord in the World Political Economy*. Princeton, NJ: Princeton University Press.

Kim, Jee-Seon and Peter M. Steiner. 2015. Multilevel Propensity Score Methods for Estimating Causal Effects: A Latent Class Modeling Strategy. In *Quantitative Psychology Research*. Springer pp. 293–306.

King, Gary and Langche Zeng. 2006. "The Dangers of Extreme Counterfactuals." *Political Analysis* 14(2):131–159.

König, Thomas and Lars Mäder. 2014. "The Strategic Nature of Compliance: An Empirical Evaluation of Law Implementation in the Central Monitoring System of the European Union." *American Journal of Political Science* 58(1):246–263.

Koremenos, Barbara, Charles Lipson and Duncan Snidal. 2001. "The Rational Design of International Institutions." *International Organization* 55(4):761–799.

Krugman, Paul. 1980. "Scale Economies, Product Differentiation, and the Pattern of Trade." *American Economic Review* 70(5):950–959.

Larsson, Olof and Daniel Naurin. 2016. "Judicial Independence and Political Uncertainty: How the Risk of Override Affects the Court of Justice of the EU." *International Organization* pp. 377–408.

Larsson, Olof, Daniel Naurin, Mattias Derlén and Johan Lindholm. 2017. "Speaking Law to Power: the Strategic use of Precedent of the Court of Justice of the European Union." *Comparative Political Studies* 50(7):879–907.

Lupu, Yonatan and Erik Voeten. 2012. "Precedent in International Courts: A Network Analysis of Case Citations by the European Court of Human Rights." *British Journal of Political Science* 42(2):413–439.

Lupu, Yonatan and James H. Fowler. 2013. "Strategic Citations to Precedent on the U.S. Supreme Court." *Journal of Legal Studies* 42(1):151–186.

Luttmer, Erzo G. J. 2007. "Selection, Growth, and the Size Distribution of Firms." *The Quarterly Journal of Economics* 122(3):1103–1144.

Mansfield, Edward D. and Diana C. Mutz. 2009. "Support for Free Trade: Self-Interest, Sociotropic Politics, and Out-Group Anxiety." *International Organization* 63(3):425.

Martinsen, Dorte Sindbjerg. 2015. "Judicial Influence on Policy Outputs? The Political Constraints of Legal Integration in the European Union." *Comparative Political Studies* 48(12):1622–1660.

Mayda, Anna Maria. 2006. "Who is Against Immigration? A Cross-Country Investigation of Individual Attitudes Toward Immigrants." *The Review of Economics and Statistics* 88(3):510–530.

Mayda, Anna Maria. 2008. "Why Are People More Pro-Trade than Pro-Migration?" *Economics Letters* 101(3):160–163.

Mayda, Anna Maria and Dani Rodrik. 2005. "Why Are Some People (and Countries) More Protectionist Than Others?" *European Economic Review* 49(6):1393–1430.

Mbaye, Heather A. D. 2001. "Why National States Comply with Supranational Law." *European Union Politics* 2(3):259–281.

Melitz, Marc J. 2003. "The Impact of Trade on Intra-Industry Reallocations and Aggregate Industry Productivity." *Econometrica* 71(6):1695–1725.

Melitz, Marc J. and Giancarlo Ottaviano. 2008. "Market Size, Trade, and Productivity." *Review of Economic Studies* 75:295–316.

Melitz, Marc J. and Stephen J. Redding. 2014. Heterogeneous firms and trade. In *Handbook of International Economics*. Vol. 4 Elsevier pp. 1–54.

O'Rourke, Kevin and Alan Taylor. 2006. "Democracy and Protectionism.".

Pandya, Sonal. 2010. "Labor Markets and the Demand for Foreign Direct Investment." *International Organization* 64(3):389–409.

Panke, Diana. 2010. "Good Instructions in No Time? Domestic Coordination of EU Policies in 19 Small States." *West European Politics* 33(4):770–790.

Pollack, Mark A. 2003. *The Engines of European Integration: Delegation, Agency and Agenda Setting in the European Union*. Oxford, England: Oxford University Press.

Rosenbaum, Paul R. and Donald B. Rubin. 1983. "The Central Role of the Propensity Score in Observational Studies for Causal Effects." *Biometrika* 70(1):41–55.

Rosendorff, B. Peter. 2005. "Stability and Rigidity: Politics and Design of the WTO's Dispute Settlement Procedure." *American Political Science Review* 99(3):389–400.

Rosendorff, Peter B. and Helen V. Milner. 2001. "The Optimal Design of International Trade Institutions: Uncertainty and Escape." *International Organization* 55(4):829–857.

Rubin, Donald B. 2004. *Multiple Imputation for Nonresponse in Surveys*. Vol. 81 John Wiley & Sons.

Scheve, Kenneth F. and Matthew J. Slaughter. 2001. "What Determines Individual Trade-Policy Preferences?" *Journal of International Economics* 54(2):267–292.

Sedelmeier, Ulrich. 2008. "After Conditionality: Post-Accession Compliance with EU Law in East Central Europe." *Journal of European Public Policy* 15(6):806–825.

Simmons, Beth A. 2009. *Mobilizing for Human Rights: International Law in Domestic Politics*. New York: Cambridge University Press.

Slepcevic, Reinhard. 2009. "The Judicial Enforcement of EU Law through National Courts: Possibilities and Limits." *Journal of European Public Policy* 16(3):378–394.

Staton, Jeffrey K. 2010. *Constitutional Review and the Selective Promotion of Case Results*. New York: Cambridge University Press.

Staton, Jeffrey K. and Will H. Moore. 2011. "Judicial Power in Domestic and International Politics." *International Organization* 65(3):553–587.

Steunenberg, Bernard and Mark Rhinard. 2010. "The Transposition of European Law in EU Member States: Between Process and Politics." *European Political Science Review* 2(3):495–520.

Stone Sweet, Alex and Thomas L. Brunell. 1998. "Constructing a Supranational Constitution: Dispute Resolution and Governance in the European Community." *American Political Science Review* 92(1):63–81.

Thomson, Robert, René Torenvlied and Javier Arregui. 2007. "The Paradox of Compliance: Infringements and Delays in Transposing European Union Directives." *British Journal of Political Science* 37(4):685–709.

Urbatsch, Robert. 2013. "A Referendum on Trade Theory: Voting on Free Trade in Costa Rica." *International Organization* 67(1):197–214.

Vanberg, Georg. 2001. "Legislative-Judicial Relations: A Game-Theoretic Approach to Constitutional Review." *American Journal of Political Science* 45(2):346–361.

Vanberg, Georg. 2005. *The Politics of Constitutional Review in Germany*. New York: Cambridge University Press.

Vanberg, Georg. 2015. "Constitutional Courts in Comparative Perspective: A Theoretical Assessment." *Annual Review of Political Science* 18:167–185.

Appendix A:

Appendix for Paper 1

A.1 Proof of Lemma 1

The implementing actor chooses p to maximize its utility:

$$
\begin{aligned}
u_A(p) &= -(p - p_i)^2 \\
\frac{\partial u_A(p)}{\partial p} &= -2(p - p_i) = 0 \\
p^* &= p_i.
\end{aligned}
$$

A.2 Proof of Lemma 2

The implementing actor's utility in equilibrium is $u_A(p^*) = -(p^* - p_i)^2 = -(p_i - p_i)^2 = 0$. The absolute difference in its utility between choosing $p_1 = p_m$ and choosing $p^* = p_i$ in the counterfactual is:

$$
\begin{aligned}
c^* &= \big|\, u_A(p_m) - u_A(p^*) \,\big| \\
&= \big|\, -(p_m - p_i)^2 \,\big| \\
&= p_i^2.
\end{aligned}
$$

A.3 Proof of Proposition 1

The plaintiff plays a cut-point strategy. It brings a case when its utility for litigating is greater than its utility for not litigating:

$$u_L(\textit{litigate}) > U_L(\neg\, \textit{litigate})$$

$$h(c^*)(-k) + (1 - h(c^*))(-(p_0 - p_m)^2 - k) > -(p_0 - p_m)^2$$

$$k < p_0^2 h(p_i^2)$$

$$k^* = p_0^2 h(p_i^2)$$

$$k < k^*.$$

Thus, the plaintiff brings a case when the cost of litigating is sufficiently low, $k < k^*$. The probability that the plaintiff brings a case in equilibrium is the probability that any given draw of k is less than the cut-point k^*, which is given by the cumulative distribution function (CDF) for k evaluated at the cut-point: $\Pr(k < k^*) = F(k^*)$.

A.4 Proof of Proposition 2

In equilibrium, the implementing actor maximizes its expected utility, which is the probability of litigation, $F(k^*)$, times its expected utility for litigation plus the probability of no litigation, 1 - $F(k^*)$, times its utility for its chosen *ex ante* policy, $u_A(p_0)$. The expected utility of litigation is the probability of *ex post* compliance, $h(c^*)$, times its utility for implementing the *de jure* policy, $u_A(p_m)$, plus the probability of *ex post* non-compliance, 1 - $h(c^*)$, times its utility for implementing its chosen *ex ante* policy, $u_A(p_0)$.

$$\mathbb{E}\big[u_A(p_0)\big] = F(k^*)\big[h(c^*)u_A(p_m) + (1 - h(c^*))u_A(p_0)\big] + (1 - F(k^*))u_A(p_0)$$

$$= p_0 h(p_i^2)(p_0 - 2p_i)F(p_0^2 h(p_i^2)) - (p_0 - p_i)^2. \tag{A.1}$$

Equation (A.1) is a convex combination of globally concave functions (quadratic loss functions). Because a convex combination is a linear combination, and because the linear combination of concave functions is itself a concave function, Equation (A.1) is a globally concave function. Thus, there is a unique interior solution. The first order condition is necessary and sufficient for a unique global maximum.

Without assigning functional forms to the functions F and h, there is not a closed-form solution for p_0^*. The implicit solution is the first order condition for Equation (A.1):

$$\frac{\partial}{\partial p_0}\mathbb{E}\big[u_A(p_0)\big] = 0$$

$$\frac{\partial}{\partial p_0}\Big[p_0 h(p_i^2)(p_0 - 2p_i)F(p_0^2 h(p_i^2)) - (p_0 - p_i)^2\Big] = 0$$

$$2p_0^2 h(p_i^2)^2(p_0 - 2p_i)F'(p_0^2 h(p_i^2)) + 2(p_0 - p_i)\Big[h(p_i^2)F(p_0^2 h(p_i^2)) - 1\Big] = 0. \tag{A.2}$$

A.5 Proof of Proposition 3

According to Proposition 2, the implementing actor's expected utility function is globally concave, which guarantees the existence of a unique optimal *ex ante* policy, p_0^*. Therefore, it must be that $p_0^* \geq p^*$ or $p_0^* < p^*$, where $p^* = p_i$ (see Lemma 1). Without loss of generality, I assume $p_i > p_m$ and I normalize $p_m = 0$. This implies $p^* > p_m$.

Suppose the implementing actor chooses $p_0 = p^* = p_i$. As shown in Equation (A.1), if there is litigation, the implementing actor's utility is a convex combination of $u_A(p_m) < 0$ and $u_A(p_0) = u_A(p^*) = 0$. If there is not litigation, the implementing actor's utility is $u_A(p_0) = 0$. Thus, $u_A(p_m) < u_A(p_0)$, and the implementing actor has an incentive to reduce the probability of litigation.

The probability of litigation is $\Pr(k < k^*) = F(k^*)$, where $k^* = p_0^2 h(p_i^2)$ is strictly increasing in p_0. Thus, if the implementing actor deviates from $p_0 = p^*$, it will choose an *ex ante* policy below the policy it chooses in the counterfactual, $p_0 < p^*$.

A.6 Computational Proofs of Results 1–5

I derive comparative statics using a combination of analytical and computational approaches. To sign the effect of a change in preference divergence (p_i) on the optimal *ex ante* policy choice (p_0), which represents the judicial impact of the court on policy implementation (Result 4), I take the total derivative Equation (A.2) with respect to preference divergence (p_i), and solve in terms of $\frac{\partial p_0}{\partial p_i}$. The total derivative is:

$$\begin{aligned}
0 = 2p_0^2 h(p_i^2)&\Big[p_0^2 p_i h(p_i^2)(p_0 - 2p_i)h'(p_i^2)F''(p_0^2 h(p_i^2)) \\
&- \Big(p_i(5p_i - 3p_0)h'(p_i^2) + h(p_i^2)\Big)F'(p_0^2 h(p_i^2))\Big] \\
&+ \frac{\partial p_0}{\partial p_i}\Big[h(p_i^2)\Big(p_0 h(p_i^2)\big[2p_0^2 h(p_i^2)(p_0 - 2p_i)F''(p_0^2 h(p_i^2)) \\
&+ (5p_0 - 6p_i)F'(p_0^2 h(p_i^2))\big] + F(p_0^2 h(p_i^2))\Big) - 1\Big] + 1 \\
&- F(p_0^2 h(p_i^2))\Big[2p_i(p_i - p_0)h'(p_i^2) + h(p_i^2)\Big].
\end{aligned}$$

We can then solve in terms of the implicit derivative of p_0^* with respect to p_i:

$$\frac{\partial p_0^*}{\partial p_i} = \frac{g_1(p_0, p_i)}{g_2(p_0, p_i)}, \qquad \text{where}$$

$$\begin{aligned}
g_1(p_0, p_i) = 2p_0^2 h(p_i^2)&\Big[\Big(p_i(5p_i - 3p_0)h'(p_i^2) + h(p_i^2)\Big)F'(p_0^2 h(p_i^2)) \\
&- p_0^2 p_i h(p_i^2)(p_0 - 2p_i)h'(p_i^2)F''(p_0^2 h(p_i^2))\Big] \\
&+ F(p_0^2 h(p_i^2))\Big[2p_i(p_i - p_0)h'(p_i^2) + h(p_i^2)\Big] - 1, \qquad \text{and}
\end{aligned}$$

$$g_2(p_0, p_i) = h(p_i^2)\Big[p_0 h(p_i^2)\Big(2p_0^2 h(p_i^2)(p_0 - 2p_i)F''(p_0^2 h(p_i^2)) \\ + (5p_0 - 6p_i)F'(p_0^2 h(p_i^2))\Big) + F(p_0^2 h(p_i^2))\Big] - 1. \tag{A.3}$$

The effect of a change in preference divergence on the effect of the court on the implementing actor's *ex ante* policy is:

$$\frac{\partial}{\partial p_i}\left[p_0^* - p^*\right] = \frac{\partial p_0^*}{\partial p_i} - \frac{\partial p^*}{\partial p_i} = \frac{\partial p_0^*}{\partial p_i} - 1. \tag{A.4}$$

To sign Equation (A.4), I choose functional forms for F and h. Assume that the functional form of F is the CDF of the exponential distribution: $F(k) = 1 - e^{-k}$, where $F'(k) = e^{-k}$ and $F''(k) = -e^{-k}$. Assume that the functional form of h is the logistic function: $h(c^*) = \frac{1}{1+e^{c^*}}$, where $h'(c^*) = \frac{-e^{c^*}}{(1+e^{c^*})^2}$. To reduce the dimensionality of the parameter space, I assume the the rate parameter of the exponential distribution and the shape parameter of the logistic function are both one. With these functional forms, the implementing actor's expected utility function is:

$$EU_G(p_0) = -\frac{p_0 e^{\left[-\frac{p_0^2}{e^{p_i^2}+1}\right]}(p_0 - 2p_i) + e^{p_i^2}(p_0 - p_i)^2 + p_i^2}{e^{p_i^2} + 1}.$$

The first order condition for this objective function is transcendental. An analytical solution for p_0^* does not exist, and it is still necessary to calculate the implicit derivative.

$$\frac{\partial}{\partial p_0} EU_G(p_0) = \frac{2\left[Xe^{\left[-\frac{p_0^2}{e^{p_i^2}+1}\right]} + e^{p_i^2}\left(e^{p_i^2} + 1\right)(p_i - p_0)\right]}{\left(e^{p_i^2} + 1\right)^2} = 0, \quad \text{where}$$

$$X = p_0\left(p_0^2 - 2p_0 p_i - 1\right) + e^{p_i^2}(p_i - p_0) + p_i.$$

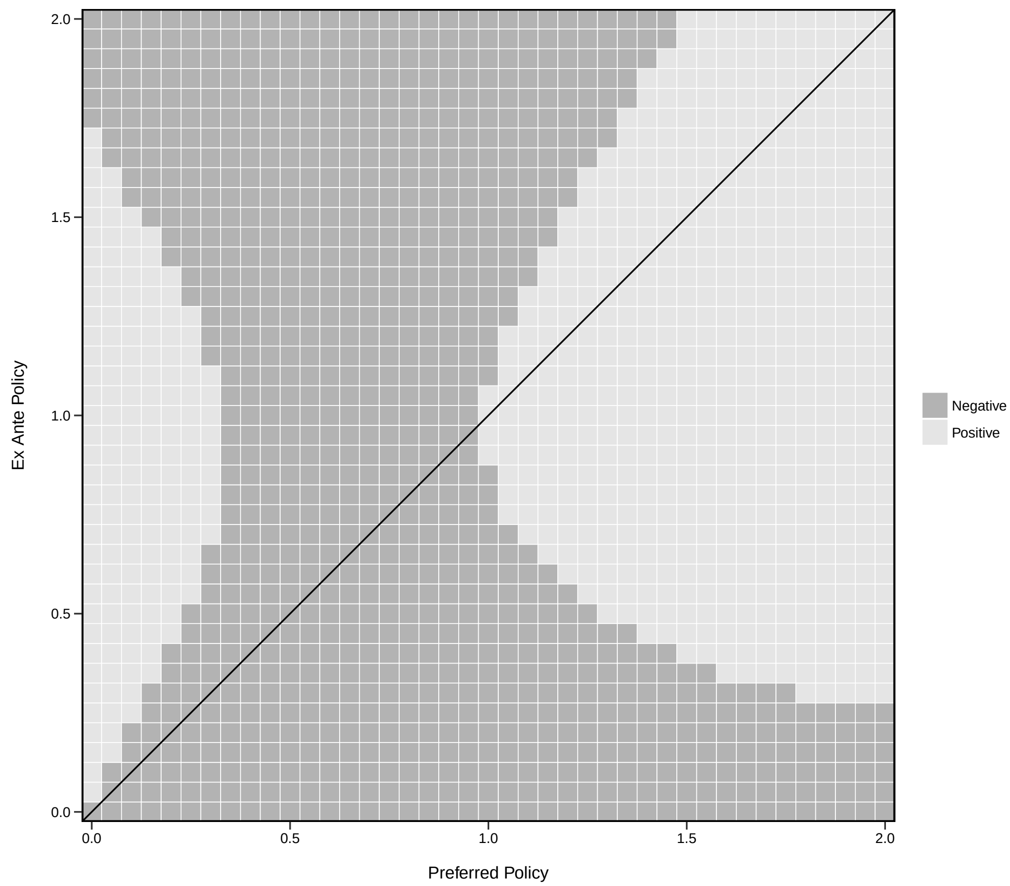

Figure A.1. Effect of Preference Divergence on the Optimal Policy

Note: This figure shows the sign of Equation (A.4) as a function of p_0 and p_i. In equilibrium, it must be that $p_0 < p_i$, which means that p_0^* must fall below the diagonal line for all p_i.

The implicit derivative in Equation (A.4) is a function of p_0 and p_i. Figure A.1 shows the sign of Equation (A.4) as a function of these two parameters. The government's optimal *ex ante* policy in equilibrium (p_0^*) must be a convex combination between its preferred policy (p_i) and the higher government's target policy ($p_m = 0$).

The government's utility function is strictly decreasing in p_0 for $p_0 > p_i$. If the implementing actor is going to implement a policy that diverges from its preferred policy ($p_0 \neq p_i$), it is going to make a concession to the higher government by implementing a policy closer to the higher government's target policy (p_m) in order to reduce the probability of litigation (which may result in *ex post* compliance). The implementing actor never has an incentive to choose a policy that is worse for it and the higher government.

In equilibrium, therefore, it must be that $p_0^* < p_i$, which means that p_0^* must fall below the diagonal line for all p_i. In the region below this diagonal line, as the government's preferred policy (p_i) increases, the implementing actor's *ex ante* policy in equilibrium (p_0^*) is decreasing, then increasing. In other words, the magnitude of the negative effect of the court on the implementing actor's *ex ante* policy is increasing, then decreasing.

I demonstrate Results 1–5 with a Monte Carlo simulation (replication code below). I numerically optimize Equation (A.2) to solve for p_0^* and other endogenous parameters (assuming the above functional forms). Note that there are no other parameters in Equation (A.2) to condition on. I use this simulation to produce Figure 1.2.

```
# solve model
solve_model <- function(pi) {

  # government utility function
  u.G <- function (p0, pi) {
    p0 * h(pi ^ 2) * (p0 - 2 * pi) * CDF((p0 ^ 2) * h(pi ^ 2)) - (p0
        - pi) ^ 2
  }

  # c.star
  c.star <- pi ^ 2

  # p0.star
  out <- optim(par = 0, fn = u.G, pi = pi, method = "Brent", lower =
      -10, upper = 10, control = list(fnscale = -1))
  p0.star <- out$par

  # k.star
  k.star <- (p0.star ^ 2) * h(pi ^ 2)

  # probability of ex post compliance
```

```
pr.ex.post.compliance <- h(c.star)

# probability of a case
pr.case <- CDF(k.star)

# ex post policy in expectation
E.p1.star <- pr.case * ((1 - pr.ex.post.compliance) * p0.star) +
    (1 - pr.case) * p0.star

# probability of successful enforcement
pr.success <- pr.case * pr.ex.post.compliance

# probability of a compliance deficit
pr.deficit <- (1 - pr.case) + pr.case * (1 - pr.ex.post.compliance
    )

# optimal concession
optimal.concession <- abs(p0.star - pi)

# ex ante compliance deficit
ex.ante.deficit <- abs(p0.star - 0)

# ex post compliance deficit in expectation
E.ex.post.deficit <- abs(E.p1.star - 0)

# solution
solution <- data.frame(pi = pi,
                       c.star = c.star,
                       pr.ex.post.compliance = pr.ex.post.
                          compliance,
                       p0.star = p0.star,
                       k.star = k.star,
                       pr.case = pr.case,
```

```
                              optimal.concession = optimal.concession,
                              E.p1.star = E.p1.star,
                              pr.success = pr.success,
                              pr.deficit = pr.deficit,
                              ex.ante.deficit = ex.ante.deficit,
                              E.ex.post.deficit = E.ex.post.deficit)

  # return
  return(solution)
}

 # solve model
solve_model_2 <- function(pi) {

  # utility function
  u.G <- function (p0, pi) {
    p0 * 0.5 * (p0 - 2 * pi) * CDF((p0 ^ 2) * 0.5) - (p0 - pi) ^ 2
  }

  # c.star
  c.star <- pi ^ 2

  # p0.star
  out <- optim(par = 0, fn = u.G, pi = pi, method = "Brent", lower =
       -10, upper = 10, control = list(fnscale = -1))
  p0.star <- out$par

  # k.star
  k.star <- (p0.star ^ 2) * 0.5

  # probability of ex post compliance
  pr.ex.post.compliance <- 0.5

```

```
# probability of a case
pr.case <- CDF(k.star)

# ex post policy in expectation
E.p1.star <- pr.case * ((1 - pr.ex.post.compliance) * p0.star) +
    (1 - pr.case) * p0.star

# probability of successful enforcement
pr.success <- pr.case * pr.ex.post.compliance

# probability of a compliance deficit
pr.deficit <- (1 - pr.case) + pr.case * (1 - pr.ex.post.compliance
    )

# optimal concession
optimal.concession <- abs(p0.star - pi)

# ex ante compliance deficit
ex.ante.deficit <- abs(p0.star - 0)

# ex post compliance deficit in expectation
E.ex.post.deficit <- abs(E.p1.star - 0)

# solution
solution <- data.frame(pi = pi,
                       c.star = c.star,
                       pr.ex.post.compliance = pr.ex.post.
                           compliance,
                       p0.star = p0.star,
                       k.star = k.star,
                       pr.case = pr.case,
                       optimal.concession = optimal.concession,
                       E.p1.star = E.p1.star,
```

```
                              pr.success = pr.success,
                              pr.deficit = pr.deficit,
                              ex.ante.deficit = ex.ante.deficit,
                              E.ex.post.deficit = E.ex.post.deficit)

  # return
  return(solution)
}

# CDF
CDF <- function(x) {
  1 - exp(- 5 * x)
}

# probability of ex post compliance
h <- function(x) {
  1 / (1 + exp(8 * (x - 0.5)))
}

################################################
# solve model
################################################

# preference for noncompliance
pi <- seq(0, 1, 0.005)

# solve model
out <- list()
for(i in 1:length(pi)) {
  out[[i]] <- solve_model(pi[i])
}

# stack observations
```

```r
out <- do.call("rbind", out)

# round estimates
out <- round(out, 5)

##################################################
# solve model (random lottery)
##################################################

# solve model
out2 <- list()
for(i in 1:length(pi)) {
  out2[[i]] <- solve_model_2(pi[i])
}

# stack observations
out2 <- do.call("rbind", out2)

# round estimates
out2 <- round(out2, 5)

##################################################
# function to make plots
##################################################

make_plot <- function(x, y, xlab = NULL, ylab = NULL, main = NULL,
    diagonal = FALSE) {
  plot <- ggplot() +
    geom_line(aes(x = x, y = y), color = "black") +
    scale_x_continuous(breaks = 0) +
    scale_y_continuous(breaks = 0) +
    xlab(xlab) +
    ylab(ylab) +
```

```
    ggtitle(main) +
    theme_bw() +
    theme(axis.text = element_text(color = "black"),
          panel.border = element_rect(size = 1, color = "black", fill = NA),
          panel.grid.major = element_blank(),
          panel.grid.minor = element_blank(),
          axis.text.x = element_text(size = 8),
          axis.text.y = element_text(size = 8),
          plot.title = element_text(size = 10, hjust = 0.5),
          axis.title.x = element_text(margin = margin(t = 3, r = 5, b = 5, l = 5), size = 10),
          axis.title.y = element_text(margin = margin(t = 5, r = 3, b = 5, l = 5), size = 10),
          axis.ticks.length = unit(5 , "pt"),
          plot.margin = margin(t = 10, r = 10, b = 10, l = 20))

  if(diagonal) {
    plot <- plot + geom_abline(slope = 1, intercept = 0, linetype = "dashed")
  }

  return(plot)
}

#################################################
# comparative statics
#################################################

# result 1(a)
# cost of compliance
result1a <- make_plot(x = out$pi, y = out$c.star, ylab = "Cost of Compliance", main = "Result 1(a)")
```

```
# result 1(b)
# probability of ex post compliance
result1b <- make_plot(x = out$pi, y = out$pr.ex.post.compliance, ylab = "Pr(Ex Post Compliance)", main = "Result 1(b)")

# Result 2
# probability of a case
result2 <- make_plot(x = out$pi, y = out$pr.case, ylab = "Pr(Litigation)", main = "Result 2")

# Result 3(a)
# ex ante policy / ex ante compliance
result3a <- make_plot(x = out$pi, y = out$p0.star, diagonal = TRUE, ylab = "Ex Ante Policy", main = "Result 3(a)")

# Result 3(b)
# expected ex post policy / expected ex post noncompliance
result3b <- make_plot(x = out$pi, y = out$E.p1.star, diagonal = TRUE, ylab = "Expected Ex Post Policy", main = "Result 3(b)")

# Result 4
# optimal concession
result4 <- make_plot(x = out$pi, y = out$optimal.concession, ylab = "Optimal Concession", main = "Result 4")

# result 5
# optimal concession (random lottery)
result5 <- make_plot(x = out2$pi, y = out2$optimal.concession, ylab = "Optimal Concession\n(Random Lottery)", main = "Result 5")

# empty plot
empty <- ggplot() + theme_void()
```

```
# align plots
plot <- ggarrange(result1a, result1b, result2, empty, result3a,
    result3b, result4, result5, ncol = 2, nrow = 4, align = "hv")
```

Appendix B:

Appendix for Paper 2

B.1 Proof of Proposition 1

The representative consumer's utility function is:

$$u_C = \left[\int_\Omega q(\omega)^{\frac{\sigma-1}{\sigma}} \, \mathrm{d}\omega \right]^{\frac{\sigma}{\sigma-1}}. \tag{B.1}$$

We can optimize $\left[u_C\right]^{\frac{\sigma-1}{\sigma}}$, which is a strictly increasing transformation of u_C, and therefore has the same optimums. The optimal quantity of each variety $q^*(\omega)$ is the solution to the following constrained optimization problem:

$$\max_{q(\omega)} \left[\int_\Omega q(\omega)^{\frac{\sigma-1}{\sigma}} \, \mathrm{d}\omega \right] \quad s.t. \quad \int_\Omega p(\omega) q(\omega) \, \mathrm{d}\omega \leq I. \tag{B.2}$$

The Lagrangian function for Equation (B.2) is:

$$\mathcal{L} = \left[\int_\Omega q(\omega)^{\frac{\sigma-1}{\sigma}} \, \mathrm{d}\omega \right] - \lambda \left[\int_\Omega p(\omega) q(\omega) \, \mathrm{d}\omega - I \right]. \tag{B.3}$$

Next, we need to find the first order condition of Equation (B.3). The consumer chooses an optimal $q(\omega)$ for each $\omega \in \Omega$. Thus, the number of first-order conditions is the mass of available varieties Ω. The first-order condition for variety ω is:

$$\begin{aligned} \frac{\partial \mathcal{L}}{\partial q(\omega)} &= 0 \\ \left[\frac{\sigma-1}{\sigma}\right] q(\omega)^{\left[\frac{\sigma-1}{\sigma}-1\right]} - \lambda p(\omega) &= 0 \\ \left[\frac{\sigma-1}{\sigma}\right] q(\omega)^{-\frac{1}{\sigma}} - \lambda p(\omega) &= 0. \end{aligned} \tag{B.4}$$

The inequality constraint is $\lambda \geq 0$. The complementary slackness condition is:

$$\lambda \left[\int_{\Omega} p(\omega) q(\omega)\, \mathrm{d}\omega - I\right] = 0.$$

The system of equations includes the first-order condition, the inequality constraint, and the complementary slackness condition:

$$\begin{aligned} \left[\frac{\sigma-1}{\sigma}\right] q(\omega)^{-\frac{1}{\sigma}} - \lambda p(\omega) &= 0 \\ \lambda &\geq 0 \\ \lambda \left[\int_{\Omega} p(\omega) q(\omega)\, \mathrm{d}\omega - I\right] &= 0. \end{aligned}$$

There are two unknowns: $q(\omega)^*$ and λ. Our intuition should be that the budget constraint is binding. Thus, the Lagrangian multiplier should be non-negative, $\lambda > 0$. Given this candidate solution, the complementary slackness conditions implies:

$$\int_{\Omega} p(\omega) q(\omega)\, \mathrm{d}\omega = I. \tag{B.5}$$

We can rearrange Equation (B.4) to solve for $q(\omega)$:

$$\begin{aligned} \left[\frac{\sigma-1}{\sigma}\right] q(\omega)^{-\frac{1}{\sigma}} - \lambda p(\omega) &= 0 \\ \left[\frac{\sigma-1}{\sigma}\right] q(\omega)^{-\frac{1}{\sigma}} &= \lambda p(\omega) \end{aligned}$$

$$
\begin{aligned}
q(\omega)^{-\frac{1}{\sigma}} &= \lambda p(\omega)\left[\frac{\sigma}{\sigma-1}\right] \\
q(\omega) &= p(\omega)^{-\sigma}\left[\frac{\lambda\sigma}{\sigma-1}\right]^{-\sigma}.
\end{aligned}
$$

The relative demand of a good ω compared to another good ω' is:

$$
\begin{aligned}
\frac{q(\omega)}{q(\omega')} &= \frac{p(\omega)^{-\sigma}\left[\frac{\lambda\sigma}{\sigma-1}\right]^{-\sigma}}{p(\omega')^{-\sigma}\left[\frac{\lambda\sigma}{\sigma-1}\right]^{-\sigma}} \\
\frac{q(\omega)}{q(\omega')} &= \frac{p(\omega)^{-\sigma}}{p(\omega')^{-\sigma}} \\
q(\omega) &= p(\omega)^{-\sigma} q(\omega') p(\omega')^{\sigma}.
\end{aligned}
$$

Next, we multiply both sides by $p(\omega)$.

$$
\begin{aligned}
q(\omega)p(\omega) &= p(\omega)^{-\sigma} p(\omega) q(\omega') p(\omega')^{\sigma} \\
q(\omega)p(\omega) &= p(\omega)^{1-\sigma} q(\omega') p(\omega')^{\sigma}.
\end{aligned}
$$

We can integrate both sides with respect to ω to aggregate:

$$
\begin{aligned}
\int_\Omega q(\omega)p(\omega)\,\mathrm{d}\omega &= \int_\Omega p(\omega)^{1-\sigma} q(\omega') p(\omega')^{\sigma}\,\mathrm{d}\omega \\
\int_\Omega q(\omega)p(\omega)\,\mathrm{d}\omega &= q(\omega') p(\omega')^{\sigma} \int_\Omega p(\omega)^{1-\sigma}\,\mathrm{d}\omega.
\end{aligned}
$$

Substituting in equation (B.5), we can find the price index:

$$
\begin{aligned}
I &= q(\omega') p(\omega')^{\sigma} \int_\Omega p(\omega)^{1-\sigma}\,\mathrm{d}\omega \\
I &= q(\omega') p(\omega')^{\sigma} (P^*)^{1-\sigma}, \qquad P^* \equiv \left[\int_\Omega p(\omega)^{1-\sigma}\,\mathrm{d}\omega\right]^{\frac{1}{1-\sigma}}.
\end{aligned} \tag{B.6}
$$

In Equation (B.6), P is the Dixit-Stiglitz price index. We can rearrange Equation (B.6) to solve for the optimal quantity of variety ω':

$$q^*(\omega) = p^*(\omega)^{-\sigma} I (P^*)^{\sigma-1}. \tag{B.7}$$

This is the quantity that maximizes Equation (B.1). The revenue that firm ω makes is the amount that the representative consumer spends on variety ω:

$$\begin{aligned} r^*(\omega) &= p^*(\omega) q^*(\omega) \\ &= p^*(\omega) p^*(\omega)^{-\sigma} I (P^*)^{\sigma-1} \\ &= p^*(\omega)^{1-\sigma} I (P^*)^{\sigma-1}. \end{aligned} \tag{B.8}$$

Firms with the same productivity ϕ behave identically. The utility of firm ϕ with respect to country $n \in N$ is given by net profit:

$$\pi(\phi) = q(\phi)\big(p(\phi) - c(\phi)\big) - f. \tag{B.9}$$

A firm's marginal cost is $c(\phi) = \frac{1}{\phi}$ in the domestic market and $c(\phi) = \frac{b}{\phi}$ in foreign markets. The firm pays fixed costs f.

In each market, firms chose a price that maximizes Equation (B.9) subject to the constraint of the equilibrium quantity:

$$\max_{p(\varphi)} \Big[q(\phi)\big(p(\phi) - c(\phi)\big) - f \Big] \quad s.t. \quad q(\phi) = p(\phi)^{-\sigma} I (P^*)^{\sigma-1}.$$

The first-order condition is:

$$\begin{aligned} \frac{\partial}{\partial p(\phi)} \Big[p(\phi)^{-\sigma} I (P^*)^{\sigma-1} \big(p(\phi) - c(\phi)\big) - f \Big] &= 0 \\ \frac{\partial}{\partial p(\phi)} \Big[p(\phi)^{1-\sigma} I (P^*)^{\sigma-1} - p(\phi)^{-\sigma} I (P^*)^{\sigma-1} c(\phi) - f \Big] &= 0. \end{aligned}$$

Next, we solve for $p(\phi)$ to get $p^*(\phi)$:

$$(1-\sigma)p(\phi)^{-\sigma}I(P^*)^{\sigma-1} + \sigma p(\phi)^{-\sigma-1}I(P^*)^{\sigma-1}c(\phi) = 0$$
$$I(P^*)^{\sigma-1}\Big[(1-\sigma)p(\phi)^{-\sigma} + \sigma p(\phi)^{-\sigma-1}c(\phi)\Big] = 0$$
$$(1-\sigma)p(\phi)^{-\sigma} + \sigma p(\phi)^{-\sigma-1}c(\phi) = 0$$
$$\begin{aligned}(1-\sigma)p(\phi)^{-\sigma} &= -\sigma p(\phi)^{-\sigma-1}c(\phi)\\ \frac{p(\phi)^{-\sigma}}{p(\phi)^{-\sigma-1}} &= -\frac{\sigma}{1-\sigma}c(\phi)\\ p^*(\phi) &= \frac{\sigma}{\sigma-1}c(\phi).\end{aligned}$$

Substituting in the marginal cost in the domestic market, the optimal price in the domestic market is:

$$p_d^*(\phi) = \frac{\sigma}{\phi(\sigma-1)}. \tag{B.10}$$

The optimal price in foreign markets is:

$$p_x^*(\phi) = \frac{b\sigma}{\phi(\sigma-1)}. \tag{B.11}$$

We can substitute Equation (B.10) into Equation (B.8) to get revenue in the domestic market:

$$r_d^*(\phi) = \left[\frac{\sigma}{\phi(\sigma-1)}\right]^{1-\sigma} I(P^*)^{\sigma-1}. \tag{B.12}$$

Similarly, we can substitute Equation (B.11) into Equation (B.8) to get revenue in foreign markets:

$$r_x^*(\phi) = \left[\frac{b\sigma}{\phi(\sigma-1)}\right]^{1-\sigma} I(P^*)^{\sigma-1}$$

$$
\begin{aligned}
&= b^{1-\sigma}\left[\frac{\sigma}{\phi(\sigma-1)}\right]^{1-\sigma} I(P^*)^{\sigma-1} \\
&= b^{1-\sigma} r_d^*(\phi). \qquad \text{(B.13)}
\end{aligned}
$$

The relative revenue of any two firms is:

$$
\begin{aligned}
\frac{r_d(\phi')}{r_d(\phi')} &= \frac{\left[\frac{\sigma}{\phi'(\sigma-1)}\right]^{1-\sigma} I(P^*)^{\sigma-1}}{\left[\frac{\sigma}{\phi''(\sigma-1)}\right]^{1-\sigma} I(P^*)^{\sigma-1}} \\
&= \left[\frac{\phi'}{\phi''}\right]^{\sigma-1}. \qquad \text{(B.14)}
\end{aligned}
$$

Substituting Equations (B.12) and (B.10) into Equation (B.9), we can solve for equilibrium profit in the domestic market:

$$
\begin{aligned}
\pi_d^*(\phi) &= q_d^*(\phi)\big(p_d^*(\phi) - c_d(\phi)\big) - f \\
&= \Big[p_d^*(\phi)^{-\sigma} I(P^*)^{\sigma-1}\Big]\Big[p_d^*(\phi) - c_d(\phi)\Big] - f \\
&= \left[\frac{\sigma}{\phi(\sigma-1)}\right]^{-\sigma} I(P^*)^{\sigma-1}\left[\frac{\sigma}{\phi(\sigma-1)} - \frac{1}{\phi}\right] - f \\
&= \left[\frac{\sigma}{\phi(\sigma-1)}\right]^{-\sigma} I(P^*)^{\sigma-1}\left[\frac{\sigma}{\phi(\sigma-1)} - \frac{(\sigma-1)}{\phi(\sigma-1)}\right] - f \\
&= \left[\frac{\sigma}{\phi(\sigma-1)}\right]^{-\sigma} I(P^*)^{\sigma-1}\left[\frac{\sigma-(\sigma-1)}{\phi(\sigma-1)}\right] - f \\
&= \left[\frac{\sigma}{\phi(\sigma-1)}\right]^{-\sigma} I(P^*)^{\sigma-1}\left[\frac{1}{\phi(\sigma-1)}\right] - f \\
&= \left[\frac{\sigma}{\phi(\sigma-1)}\right]^{-\sigma} I(P^*)^{\sigma-1}\left[\frac{\sigma}{\phi(\sigma-1)}\right]\frac{1}{\sigma} - f \\
&= \frac{1}{\sigma}\left[\frac{\sigma}{\phi(\sigma-1)}\right]^{1-\sigma} I(P^*)^{\sigma-1} - f \\
&= \frac{r_d^*(\phi)}{\sigma} - f. \qquad \text{(B.15)}
\end{aligned}
$$

Substituting Equations (B.13) and (B.11) into Equation (B.9), we can solve for equilibrium profit in foreign markets:

$$\pi_x^*(\phi) = \frac{r_x^*(\phi)}{\sigma} - f. \tag{B.16}$$

Firms produce for a market when net profit is positive:

$$\begin{aligned} \pi^*(\phi) &> 0 \\ \frac{r^*(\phi)}{\sigma} - f &> 0 \\ r^*(\phi) &> \sigma f. \end{aligned}$$

Firms make zero profits when:

$$r^*(\phi) = \sigma f. \tag{B.17}$$

Firms play a cut-point strategy. The domestic cut-point ϕ_d^* is the productivity above which firms make profits in the domestic market:

$$\begin{aligned} r_d^*(\phi) &> \sigma f \\ \left[\frac{\sigma}{\phi(\sigma-1)}\right]^{1-\sigma} I(P^*)^{\sigma-1} &> \sigma f \\ \left[\frac{\sigma}{\phi(\sigma-1)}\right]^{1-\sigma} &> \frac{\sigma f}{I(P^*)^{\sigma-1}} \\ \left[\frac{\sigma}{\phi(\sigma-1)}\right] &> \left[\frac{\sigma f}{I(P^*)^{\sigma-1}}\right]^{\frac{1}{1-\sigma}} \\ \frac{1}{\phi}\left[\frac{\sigma}{\sigma-1}\right] &> \left[\frac{\sigma f}{I(P^*)^{\sigma-1}}\right]^{\frac{1}{1-\sigma}} \\ \frac{1}{\phi} &> \left[\frac{\sigma-1}{\sigma}\right]\left[\frac{\sigma f}{I(P^*)^{\sigma-1}}\right]^{\frac{1}{1-\sigma}} \\ \phi &> \left[\frac{\sigma}{\sigma-1}\right]\left[\frac{\sigma f}{I(P^*)^{\sigma-1}}\right]^{\frac{1}{\sigma-1}} \end{aligned}$$

$$\phi > \phi_d^* \equiv \left[\frac{\sigma}{\sigma - 1}\right] \left[\frac{\sigma f}{I(P^*)^{\sigma-1}}\right]^{\frac{1}{\sigma-1}}. \tag{B.18}$$

There is also an exporting cut-point ϕ_x^*, which is the productivity at which firms make zero profit in a foreign market. We can write the exporting cut-point in terms of the domestic cut-point ϕ_d^*. Using Equation (B.14), we can compare the domestic revenue of firms at the domestic cut-point and at the exporting cut-point:

$$\begin{aligned} \frac{r_d(\phi_x^*)}{r_d(\phi_d^*)} &= \left[\frac{\phi_x^*}{\phi_d^*}\right]^{\sigma-1} \\ r_d(\phi_x^*) &= \left[\frac{\phi_x^*}{\phi_d^*}\right]^{\sigma-1} r_d(\phi_d^*). \end{aligned} \tag{B.19}$$

Then, using Equation (B.13), we can express foreign revenue at the exporting cut-point $r_x(\phi_x^*)$ in terms of domestic revenue at the exporting cut-point $r_d(\phi_x^*)$:

$$r_x(\phi_x^*) = b^{1-\sigma} r_d(\phi_x^*). \tag{B.20}$$

We can now use Equations (B.19) and (B.20) to express foreign revenue at the exporting cut-point $r_x(\phi_x^*)$ in terms of domestic revenue at the domestic cut-point $r_d(\phi_d^*)$:

$$\begin{aligned} r_x(\phi_x^*) &= b^{1-\sigma} r_d(\phi_x^*) \\ r_x(\phi_x^*) &= b^{1-\sigma} \left[\frac{\phi_x^*}{\phi_d^*}\right]^{\sigma-1} r_d(\phi_d^*). \end{aligned} \tag{B.21}$$

Then, we can substitute Equation (B.17) into Equation (B.21) to solve for the exporting cut-point ϕ_x^* in terms of the domestic cut-point ϕ_d^*:

$$\begin{aligned} r_x(\phi_x^*) &= b^{1-\sigma} \left[\frac{\phi_x^*}{\phi_d^*}\right]^{\sigma-1} r_d(\phi_d^*) \\ \sigma f &= b^{1-\sigma} \left[\frac{\phi_x^*}{\phi_d^*}\right]^{\sigma-1} \sigma f \\ 1 &= b^{1-\sigma} (\phi_d^*)^{1-\sigma} (\phi_x^*)^{\sigma-1} \end{aligned}$$

$$(\phi_x^*)^{1-\sigma} = b^{1-\sigma}(\phi_d^*)^{1-\sigma}$$
$$\phi_x^* = b\phi_d^*. \tag{B.22}$$

Since $b \geq 1$, the exporting cut-point ϕ_x^* is always greater than the domestic cut-point ϕ_d^*. This means that only a fraction of the firms that product for the domestic market will export in equilibrium.

To solve for the domestic cut-point ϕ_d^* (see Equation B.18), we need to solve for the equilibrium price index P^*. These are the only remaining unknowns.

Assume that productivity ϕ is Pareto distributed with probability density function $g(\phi)$ and cumulative distribution function $G(\phi)$:

$$g(\phi, \theta) = \frac{\theta}{\phi^{\theta+1}}$$
$$G(\phi, \theta) = 1 - \phi^{-\theta}.$$

The shape parameter θ indicates the homogeneity of firms. Higher values of θ indicate a more homogenous sector. We must assume $\theta > \sigma - 1$ for the average productivity of firms in equilibrium to be finite. Otherwise, there is an integral in the expression for that quantity that will be divergent.

The probably that a firm supplies the domestic market is $1 - G(\phi_d^*)$. The mass of firms that supply the domestic market is:

$$\begin{aligned} M_d^* &= (1 - G(\phi_d^*))M \\ &= (1 - (1 - (\phi_d^*)^{-\theta}))M \\ &= (\phi_d^*)^{-\theta} M. \end{aligned} \tag{B.23}$$

The probability that a firm exports is $1 - G(\phi_x^*)$ and the mass of exporting firms is:

$$M_x^* = (\phi_x^*)^{-\theta} M. \tag{B.24}$$

The price index is a weighted average of prices for all goods available in a market, which includes goods supplied by domestic firms and foreign firms:

$$\begin{aligned}
P^* &\equiv \left[\int_\Omega p(\omega)^{1-\sigma}\,\mathrm{d}\omega\right]^{\frac{1}{1-\sigma}} \\
P^* &= \left[\int_{\phi_d^*}^\infty p_d(\phi)^{1-\sigma} M_d^* \frac{g(\phi)}{1-G(\phi_d^*)}\,\mathrm{d}\phi + n\int_{\phi_x^*}^\infty p_x(\phi)^{1-\sigma} M_x^* \frac{g(\phi)}{1-G(\phi_x^*)}\,\mathrm{d}\phi\right]^{\frac{1}{1-\sigma}} \\
&= \left[M_d^* \int_{\phi_d^*}^\infty \left[\frac{\sigma}{\phi(\sigma-1)}\right]^{1-\sigma} \frac{g(\phi)}{1-G(\phi_d^*)}\,\mathrm{d}\phi + nM_x^* \int_{\phi_x^*}^\infty \left[\frac{b\sigma}{\phi(\sigma-1)}\right]^{1-\sigma} \frac{g(\phi)}{1-G(\phi_x^*)}\,\mathrm{d}\phi\right]^{\frac{1}{1-\sigma}} \\
&= \left[\frac{\sigma}{\sigma-1}\right]\left[M_d^* \int_{\phi_d^*}^\infty \phi^{\sigma-1}\frac{g(\phi)}{1-G(\phi_d^*)}\,\mathrm{d}\phi + nM_x^*(b^{1-\sigma})\int_{\phi_x^*}^\infty \phi^{\sigma-1}\frac{g(\phi)}{1-G(\phi_x^*)}\,\mathrm{d}\phi\right]^{\frac{1}{1-\sigma}} \\
&= \left[\frac{\sigma}{\sigma-1}\right]\left[M_d^*\tilde{\phi}_d^* + nM_x^*(b^{1-\sigma})\tilde{\phi}_x^*\right]^{\frac{1}{1-\sigma}}, \qquad \tilde{\phi}^* \equiv \int_{\phi^*}^\infty \phi^{\sigma-1}\frac{g(\phi)}{1-G(\phi^*)}\,\mathrm{d}\phi.
\end{aligned} \tag{B.25}$$

The average productivity of firms supplying the domestic market is:

$$\begin{aligned}
\tilde{\phi}_d^* &= \int_{\phi_d^*}^\infty \phi^{\sigma-1}\frac{g(\phi)}{1-G(\phi_d^*)}\,\mathrm{d}\phi \\
&= \int_{\phi_d^*}^\infty \phi^{\sigma-1}\frac{\frac{\theta}{\phi^{\theta+1}}}{1-(1-(\phi_d^*)^{-\theta})}\,\mathrm{d}\phi \\
&= \int_{\phi_d^*}^\infty \phi^{\sigma-1}\frac{\theta}{\phi^{\theta+1}}(\phi_d^*)^\theta\,\mathrm{d}\phi \\
&= \theta(\phi_d^*)^\theta\left[\int_{\phi_d^*}^\infty \frac{\phi^{\sigma-1}}{\phi^{\theta+1}}\,\mathrm{d}\phi\right] \\
&= \theta(\phi_d^*)^\theta\left[\int_{\phi_d^*}^\infty \phi^{\sigma-\theta-2}\,\mathrm{d}\phi\right] \\
&= \theta(\phi_d^*)^\theta\left[\frac{(\phi_d^*)^{\sigma-\theta-1}}{1-\sigma+\theta}\right] \\
&= \frac{\theta(\phi_d^*)^{\sigma-1}}{1-\sigma+\theta}.
\end{aligned} \tag{B.26}$$

Note that the average productivity of firms supplying the domestic market $\tilde{\phi}_d^*$ only depends on the equilibrium domestic cut-point ϕ_d^*.

Similarly, the average productivity of firms supplying foreign markets is:

$$\tilde{\phi}_x^* = \frac{\theta(\phi_x^*)^{\sigma-1}}{1-\sigma+\theta}. \tag{B.27}$$

Substituting Equations (B.22), (B.23), (B.24), (B.26), and (B.27) into Equation (B.25), we can write the price index in terms of the domestic cut-point ϕ_d^*:

$$\begin{aligned}
P^* &= \left[\frac{\sigma}{\sigma-1}\right]\left[M_d^*\tilde{\phi}_d^* + nM_x^*(b^{1-\sigma})\tilde{\phi}_x^*\right]^{\frac{1}{1-\sigma}} \\
&= \left[\frac{\sigma}{\sigma-1}\right]\left[(\phi_d^*)^{-\theta}M\tilde{\phi}_d^* + n(\phi_x^*)^{-\theta}M(b^{1-\sigma})\tilde{\phi}_x)\right]^{\frac{1}{1-\sigma}} \\
&= \left[\frac{\sigma}{\sigma-1}\right]\left[(\phi_d^*)^{-\theta}M\tilde{\phi}_d^* + n(b\phi_d^*)^{-\theta}M(b^{1-\sigma})\tilde{\phi}_x^*\right]^{\frac{1}{1-\sigma}} \\
&= \left[\frac{\sigma}{\sigma-1}\right]\left[(\phi_d^*)^{-\theta}M\left[\frac{\theta(\phi_d^*)^{\sigma-1}}{1-\sigma+\theta}\right] + n(b\phi_d^*)^{-\theta}M(b^{1-\sigma})\left[\frac{\theta(\phi_x^*)^{\sigma-1}}{1-\sigma+\theta}\right]\right]^{\frac{1}{1-\sigma}} \\
&= \left[\frac{\sigma}{\sigma-1}\right]\left[(\phi_d^*)^{-\theta}M\left[\frac{\theta(\phi_d^*)^{\sigma-1}}{1-\sigma+\theta}\right] + n(b\phi_d^*)^{-\theta}M(b^{1-\sigma})\left[\frac{\theta(b\phi_d^*)^{\sigma-1}}{1-\sigma+\theta}\right]\right]^{\frac{1}{1-\sigma}} \\
&= \left[\frac{\sigma}{\sigma-1}\right]\left[(\phi_d^*)^{-\theta}M\left[\frac{\theta(\phi_d^*)^{\sigma-1}}{1-\sigma+\theta}\right] + n(b\phi_d^*)^{-\theta}M(b^{1-\sigma})\left[\frac{\theta(b\phi_d^*)^{\sigma-1}}{1-\sigma+\theta}\right]\right]^{\frac{1}{1-\sigma}} \\
&= \left[\frac{\sigma}{\sigma-1}\right]\left[\frac{(\phi_d^*)^{-\theta}M\theta(\phi_d^*)^{\sigma-1} + n(b\phi_d^*)^{-\theta}M(b^{1-\sigma})\theta(b\phi_d^*)^{\sigma-1}}{1-\sigma+\theta}\right]^{\frac{1}{1-\sigma}} \\
&= \left[\frac{\sigma}{\sigma-1}\right]\left[\frac{(\phi_d^*)^{\sigma-\theta-1}M\theta + (\phi_d^*)^{\sigma-\theta-1}nM\theta b^{-\theta}b^{1-\sigma}b^{\sigma-1}}{1-\sigma+\theta}\right]^{\frac{1}{1-\sigma}} \\
&= \left[\frac{\sigma}{\sigma-1}\right]\left[\frac{(\phi_d^*)^{\sigma-\theta-1}M\theta\left[1+nb^{-\theta}\right]}{1-\sigma+\theta}\right]^{\frac{1}{1-\sigma}} \\
&= (\phi_d^*)^{\frac{\sigma-\theta-1}{1-\sigma}}\left[\frac{\sigma}{\sigma-1}\right]\left[\frac{M\theta\left[1+nb^{-\theta}\right]}{1-\sigma+\theta}\right]^{\frac{1}{1-\sigma}}.
\end{aligned} \tag{B.28}$$

Now we can plug the equilibrium price index P^* into the equilibrium domestic cut-point ϕ_d^* to get an equation with one unknown:

$$\begin{aligned}
\phi_d^* &= \left[\frac{\sigma}{\sigma-1}\right]\left[\frac{\sigma f}{I(P^*)^{\sigma-1}}\right]^{\frac{1}{\sigma-1}} \\
&= \left[\frac{\sigma}{\sigma-1}\right]\left[\sigma f I^{-1}\right]^{\frac{1}{\sigma-1}}\frac{1}{P^*} \\
&= \left[\frac{\sigma}{\sigma-1}\right]\left[\sigma f I^{-1}\right]^{\frac{1}{\sigma-1}}(\phi_d^*)^{\frac{\sigma-\theta-1}{1-\sigma}}\left[\frac{\sigma}{\sigma-1}\right]\left[\frac{M\theta\left[1+nb^{-\theta}\right]}{1-\sigma+\theta}\right]^{\frac{1}{1-\sigma}} \\
&= (\phi_d^*)^{\frac{\sigma-\theta-1}{\sigma-1}}\left[\frac{\sigma f M\theta\left[1+nb^{-\theta}\right]}{I(1-\sigma+\theta)}\right]^{\frac{1}{\sigma-1}} \\
\phi_d^* &= (\phi_d^*)^{\frac{\sigma-\theta-1}{\sigma-1}}X^{\frac{1}{\sigma-1}}, \qquad X \equiv \left[\frac{\sigma f M\theta\left[1+nb^{-\theta}\right]}{I(1-\sigma+\theta)}\right].
\end{aligned}$$

With ϕ_d^* isolated, we can solve for a closed-form solution:

$$\begin{aligned}
\phi_d^* &= (\phi_d^*)^{\frac{\sigma-\theta-1}{\sigma-1}}X^{\frac{1}{\sigma-1}} \\
\phi_d^* - (\phi_d^*)^{\frac{\sigma-\theta-1}{\sigma-1}} &= X^{\frac{1}{\sigma-1}} \\
(\phi_d^*)^{1-\frac{\sigma-\theta-1}{\sigma-1}} &= X^{\frac{1}{\sigma-1}} \\
(\phi_d^*)^{\frac{\sigma-1}{\sigma-1}-\frac{\sigma-\theta-1}{\sigma-1}} &= X^{\frac{1}{\sigma-1}} \\
(\phi_d^*)^{\frac{\theta}{\sigma-1}} &= X^{\frac{1}{\sigma-1}} \\
\phi_d^* &= \left[X^{\frac{1}{\sigma-1}}\right]^{\frac{\sigma-1}{\theta}} \\
\phi_d^* &= X^{\frac{1}{\theta}} \\
\phi_d^* &= \left[\frac{\sigma f M\theta\left[1+nb^{-\theta}\right]}{I(1-\sigma+\theta)}\right]^{\frac{1}{\theta}}.
\end{aligned} \tag{B.29}$$

Substituting Equation (B.29) into Equation (B.28), we can get a closed-form solution for the equilibrium price index P^*:

$$\begin{aligned}
P^* &= (\phi_d^*)^{\frac{\sigma-\theta-1}{1-\sigma}} \left[\frac{\sigma}{\sigma-1}\right] \left[\frac{M\theta\Big[1+nb^{-\theta}\Big]}{1-\sigma+\theta}\right]^{\frac{1}{1-\sigma}} \\
&= \left[\left[\frac{\sigma f M\theta\Big[1+nb^{-\theta}\Big]}{I(1-\sigma+\theta)}\right]^{\frac{1}{\theta}}\right]^{\frac{\sigma-\theta-1}{1-\sigma}} \left[\frac{\sigma}{\sigma-1}\right] \left[\frac{M\theta\Big[1+nb^{-\theta}\Big]}{1-\sigma+\theta}\right]^{\frac{1}{1-\sigma}} \\
&= \left[\frac{\sigma}{\sigma-1}\right] \left[\frac{\sigma f M\theta\Big[1+nb^{-\theta}\Big]}{I(1-\sigma+\theta)}\right]^{\frac{\sigma-\theta-1}{\theta(1-\sigma)}} \left[\frac{M\theta\Big[1+nb^{-\theta}\Big]}{1-\sigma+\theta}\right]^{\frac{1}{1-\sigma}} \\
&= \left[\frac{\sigma}{\sigma-1}\right] \left[\frac{\sigma f}{I}\right]^{\frac{\sigma-\theta-1}{\theta(1-\sigma)}} \left[\frac{M\theta\Big[1+nb^{-\theta}\Big]}{1-\sigma+\theta}\right]^{-\frac{1}{\theta}} . \qquad \text{(B.30)}
\end{aligned}$$

In sum, equilibrium behavior is given by:

$$\text{consumption} \quad \begin{cases} q_d^*(\phi) = \left[\frac{\sigma}{\phi(\sigma-1)}\right]^{-\sigma} I(P^*)^{\sigma-1} \\[2ex] q_x^*(\phi) = \left[\frac{b\sigma}{\phi(\sigma-1)}\right]^{-\sigma} I(P^*)^{\sigma-1} \end{cases}$$

$$\text{pricing} \quad \begin{cases} p_d^*(\phi) = \frac{\sigma}{\phi(\sigma-1)} \\[2ex] p_x^*(\phi) = \frac{b\sigma}{\phi(\sigma-1)} \end{cases}$$

$$\text{production} \quad \begin{cases} \phi_d^* = \left[\frac{\sigma f M\theta\left[1+nb^{-\theta}\right]}{I(1-\sigma+\theta)}\right]^{\frac{1}{\theta}} \\[2ex] \phi_x^* = b\left[\frac{\sigma f M\theta\left[1+nb^{-\theta}\right]}{I(1-\sigma+\theta)}\right]^{\frac{1}{\theta}} . \end{cases}$$

The other equilibrium quantities of interest are:

$$M_d^* = (\phi_d^*)^{-\theta} M = M \left[\frac{\sigma f M \theta \left[1 + nb^{-\theta}\right]}{I(1-\sigma+\theta)} \right]^{-1}$$

$$P^* = \left[\frac{\sigma}{\sigma-1}\right] \left[\frac{\sigma f}{I}\right]^{\frac{\sigma-\theta-1}{\theta(1-\sigma)}} \left[\frac{M\theta\left[1+nb^{-\theta}\right]}{1-\sigma+\theta} \right]^{-\frac{1}{\theta}}$$

$$W^* = \frac{I}{P^*} = I \left[\frac{\sigma-1}{\sigma}\right] \left[\frac{\sigma f}{I}\right]^{-\frac{\sigma-\theta-1}{\theta(1-\sigma)}} \left[\frac{M\theta\left[1+nb^{-\theta}\right]}{1-\sigma+\theta} \right]^{\frac{1}{\theta}}.$$

B.2 Proof of Proposition 2

The government's utility is:

$$\begin{aligned} u_g(b) &= -w\phi_d^* - (1-w)\phi_x^* \\ &= -w\phi_d^* - (1-w)b\phi_d^* \\ &= -w \left[\frac{\sigma f M\theta\left[1+nb^{-\theta}\right]}{I(1-\sigma+\theta)} \right]^{\frac{1}{\theta}} - (1-w)b \left[\frac{\sigma f M\theta\left[1+nb^{-\theta}\right]}{I(1-\sigma+\theta)} \right]^{\frac{1}{\theta}} \\ &= -w\left[1+nb^{-\theta}\right]^{\frac{1}{\theta}} \left[\frac{\sigma f M\theta}{I(1-\sigma+\theta)}\right]^{\frac{1}{\theta}} - (1-w)b\left[1+nb^{-\theta}\right]^{\frac{1}{\theta}} \left[\frac{\sigma f M\theta}{I(1-\sigma+\theta)}\right]^{\frac{1}{\theta}} \\ &= -w\left[1+nb^{-\theta}\right]^{\frac{1}{\theta}} A - (1-w)b\left[1+nb^{-\theta}\right]^{\frac{1}{\theta}} A, \quad A \equiv \left[\frac{\sigma f M\theta}{I(1-\sigma+\theta)}\right]^{\frac{1}{\theta}}. \end{aligned}$$

The first-order condition is:

$$\begin{aligned} \frac{\partial}{\partial b} - w\left[1+nb^{-\theta}\right]^{\frac{1}{\theta}} A - (1-w)b\left[1+nb^{-\theta}\right]^{\frac{1}{\theta}} A = 0 \\ \frac{A(1+b^{-\theta}n)^{\frac{1}{\theta}}\left(b^{1+\theta}(w-1)+nw\right)}{b(b^{\theta}+n)} = 0. \end{aligned}$$

Now we can solve for b to get $\tilde{b}^*$:

$$
\begin{aligned}
\frac{A(1+b^{-\theta}n)^{\frac{1}{\theta}}\big(b^{1+\theta}(w-1)+nw\big)}{b(b^{\theta}+n)} &= 0 \\
A(1+b^{-\theta}n)^{\frac{1}{\theta}}\big(b^{1+\theta}(w-1)+nw\big) &= 0 \\
b^{1+\theta}(w-1)+nw &= 0 \\
b^{1+\theta}(w-1) &= -nw \\
b^{1+\theta} &= \frac{-nw}{(w-1)} \\
\tilde{b}^* &= \left[\frac{nw}{1-w}\right]^{\frac{1}{1+\theta}}.
\end{aligned}
$$

B.3 Proof of Proposition 3

A litigant brings a case when the costs are sufficiently small:

$$
\begin{aligned}
h(c^*)(-k)+\big(1-h(c^*)\big)\big(-(b_t-b_0)^2-k\big) &> -(b_t-b_0)^2 \\
-kh(c^*)-(b_t-b_0)^2-k+h(c^*)(b_t-b_0)^2+kh(c^*) &> -(b_t-b_0)^2 \\
-(b_t-b_0)^2-k+h(c^*)(b_t-b_0)^2 &> -(b_t-b_0)^2 \\
-k &> -h(c^*)(b_t-b_0)^2 \\
k &< h(c^*)(b_t-b_0)^2 \\
k &< k^* \equiv h(c^*)(b_t-b_0)^2.
\end{aligned}
$$

The probability that a litigant brings a case is $\Pr(k < k^*) = J(k^*)$.

B.4 Computational Proof of Proposition 4

Assume $b_t < \tilde{b}^*$. The expected utility of a government is:

$$E\Big[u_g(b_0)\Big] = J(k^*)\Big[h(c^*)u_g(b_t) + \big(1 - h(c^*)\big)u_g(b_0)\Big] + \big(1 - J(k^*)\big)u_g(b_0).$$

Assume the functional form of $J(k)$ is the exponential CDF:

$$J(k) = 1 - e^{-k}.$$

And assume the functional form of $h(c)$ is the logistic function:

$$h(c^*) = \frac{1}{1 + e^{c^*}}.$$

The cost of compliance in equilibrium is:

$$\begin{aligned} c^* &= \Big|\, u_g(\tilde{b}^*) - u_g(b_t)\, \Big| \\ &= \Bigg|\, \Big[- w\phi_d^*(\tilde{b}^*) - (1-w)\tilde{\phi}_x^*(\tilde{b}^*)\Big] - \Big[- w\phi_d^*(b_t) - (1-w)\phi_x^*(b_t)\Big]\, \Bigg|. \end{aligned}$$

Equilibrium *ex ante* trade barriers b_0^* maximize the expected utility of the government, $E\big[u_g(b_0)\big]$. Due to the complexity of this optimization problem, I calculate b_0^* numerically. The expected utility function is globally concave, which guarantees a unique solution (which can be a boundary solution for some combinations of parameter values). I program a numerical simulation of the economy in `R` using the analytical solutions for φ_d^* and $\tilde{b}^*$. I use this simulation to calculate a numerical solution for b_0^* and to derive comparative statics computationally. I build up the simulation in stages, starting with the economy.

```
# function to solve the economy using the analytical solution
solve.economy.analytical <- function(var) {
```

```r

  # solve for the domestic production cut-point
  var$phi.d.star <- (var$sigma * var$f * var$M * var$theta  * (1 + var$n * var$b ^ (-var$theta)) / (var$I * (1 - var$sigma + var$theta))) ^ (1 / var$theta)

  # identify boundary solutions
  var$phi.d.star <- max(var$phi.d.star, 1)

  # solve for the exporting cut-point
  var$phi.x.star <- var$b * var$phi.d.star

  # solve for the mass of firms that produce for the domestic market
  var$M.d.star <- var$phi.d.star ^ (-var$theta) * var$M

  # solve for the mass of firms that produce for foreign markets
  var$M.x.star <- var$phi.x.star ^ (-var$theta) * var$M

  # solve for the price index
  var$P.star <- (var$sigma / (var$sigma - 1)) * (var$sigma * var$f / var$I) ^ ((var$sigma - var$theta - 1) / (var$theta * (1 - var$sigma))) * ((var$M * var$theta * (1 + var$n * var$b ^ (-var$theta))) / (1 - var$sigma + var$theta)) ^ (-1 / var$theta)

  # solve for consumer welfare
  var$W.star <- 1 / var$P.star

  # return solution
  return(var)
}
```

I confirm the analytical solution for φ_d^* by comparing a numerical simulation that uses the analytical solution for φ_d^* to a numerical simulation that uses a computational solution for φ_d^*. I calculate the other endogenous parameters in the economy conditional on the numerical value of φ_d^*.

I solve the economy as a system of nonlinear equations with one unknown, φ_d^*, using the Newton method. The Newton method calculates the Jacobian of the function at each iteration. I use an implementation of the Newton method from `nleqslv`.

```
# function to solve the economy computationally
solve.economy.computational <- function(var) {

  # Pareto PDF
  g <- function(phi, theta) {
    theta / phi ^ (theta + 1)
  }

  # Pareto CDF
  G <- function(phi, theta) {
    1 - phi ^ (-theta)
  }

  # function to solve the system of equations
  solve.system <- function(var, phi.d.star) {

    # solve for the exporting cut-point
    phi.x.star <- var$b * phi.d.star

    # calculate rho
    rho <- ((var$sigma - 1) / var$sigma)

    # solve for the mass of firms
    M.d <- (1 - G(phi.d.star, var$theta)) * var$M
    M.x <- (1 - G(phi.x.star, var$theta)) * var$M
    M.t <- M.d + var$n * M.x
```

```
  # solve for average productivity
  phi.d.tilde <- ((phi.d.star ^ (var$sigma - 1) * var$theta ) / (1
      - var$sigma + var$theta)) ^ (1 / (var$sigma - 1))
  phi.x.tilde <- ((phi.x.star ^ (var$sigma - 1) * var$theta ) / (1
      - var$sigma + var$theta)) ^ (1 / (var$sigma - 1))
  phi.t.tilde <- ((1 / M.t) * (M.d * phi.d.tilde ^ (var$sigma - 1)
      + var$n * M.x * ((1 / var$b) * phi.x.tilde) ^ (var$sigma -
     1))) ^ (1 / (var$sigma - 1))

  # solve for the price index
  P <- M.t ^ (1 / (1 - var$sigma)) * (1 / (rho * phi.t.tilde))

  # solve for the domestic production cut-point
  out <- ((rho * phi.d.star) ^ (var$sigma - 1) * var$I * P ^ (var$
     sigma - 1)) / var$sigma - var$f

  # return solution
  return(out)
}

# solve for the domestic production cut-point
out <- nleqslv(x = runif(1, 1, 2), fn = solve.system, var = var,
    method = "Newton", control = list(maxit = 1000, allowSingular =
     TRUE))
var$phi.d.star <- out$x

# calculate rho
rho <- ((var$sigma - 1) / var$sigma)

# solve for the exporting cut-point
var$phi.x.star <- var$b * var$phi.d.star

```

```
  # solve for the mass of firms
  var$M.d.star <- (1 - G(var$phi.d.star, var$theta)) * var$M
  var$M.x.star <- (1 - G(var$phi.x.star, var$theta)) * var$M
  M.t <- var$M.d.star + var$n * var$M.x.star

  # solve for average productivity
  phi.d.tilde <- ((var$phi.d.star ^ (var$sigma - 1) * var$theta ) / (1 - var$sigma + var$theta)) ^ (1 / (var$sigma - 1))
  phi.x.tilde <- ((var$phi.x.star ^ (var$sigma - 1) * var$theta ) / (1 - var$sigma + var$theta)) ^ (1 / (var$sigma - 1))
  phi.t.tilde <- ((1 / M.t) * (var$M.d * phi.d.tilde ^ (var$sigma - 1) + var$n * var$M.x.star * ((1 / var$b) * phi.x.tilde) ^ (var$sigma - 1))) ^ (1 / (var$sigma - 1))

  # solve for the price index
  var$P.star <- M.t ^ (1 / (1 - var$sigma)) * (1 / (rho * phi.t.tilde))

  # solve for consumer welfare
  var$W.star <- 1 / var$P.star

  # return solution
  return(var)
}
```

Next, I program a numerical simulation of the counterfactual model (no regime) that builds on my analytical solution for $\tilde{b}^*$ and my numerical simulation of the economy. To maximize computational speed, I use the implementation of my numerical simulation that uses my analytical solution for optimal trade barriers $\tilde{b}^*$.

```
# function to solve the counterfactual model using the analytical solution
solve.counterfactual.analytical <- function(var) {
```

```

  # calculate optimal trade barriers
  var$b.tilde.star <- (-(var$w - 1) / (var$n * var$w)) ^ (1 / (-1 - var$theta))
  var$b <- var$b.tilde.star

  # solve the economy
  var <- solve.economy.analytical(var)
  var$b <- NULL

  # return solution
  return(var)
}
```

I confirm the analytical solution for $\tilde{b}^*$ using numerical optimization. I use `optim` in `R`. To improve performance, I optimize over the numerical solution to the economy that uses the analytical solution to φ_d^*.

```
# function to solve the model without a regime
solve.counterfactual.computational <- function (var) {

  # function to calculate the government's utility
  u.G <- function (b.tilde, var) {

    # set the variable costs of trade equal to the government's choice of trade barriers
    var$b <- b.tilde

    # solve the economy conditional on the variable costs of trade
    var <- solve.economy.computational(var)

    # calculate the government's utility
    u.G <- - var$w * var$phi.d.star - (1 - var$w) * var$phi.x.star
```

```

    # return the government's utility
    return (u.G)
  }

  # function to calculate the optimal trade barriers
  b.tilde.star <- function(var) {

    # optimize the government's utility function
    out <- optim(fn = u.G, var = var, par = runif(1, 0, 1), method =
        "Brent", lower = 1, upper = 10, control = list(fnscale = -1)
      )

    # save the government's equilibrium trade barriers
    var$b.tilde.star <- out$par

    # return the list of parameters
    return(var)
  }

  # calculate the government's optimal policy
  var <- b.tilde.star(var)

  # set the variable cost of trade equal to equilibrium trade
     barriers
  var$v <- var$b.tilde.star

  # plug the equilibrium trade barriers into the economy and solve
  var <- solve.economy.computational(var)

  # return the solution
  return(var)
}
```

Turning to the full model, I program a function that calculates the numerical solution for *ex ante* trade barriers in equilibrium, b_0^*. I use analytical solutions to program the expected utility function of the government as a function with one parameter, b_0. I express all other endogenous parameters as a function of b_0 and exogenous parameters.

I calculate the b_0^* using numeric optimization (`optim`). At each iteration of the optimization routine, I calculate all of the endogenous parameters of the model using analytical solutions, conditional on the current estimate of b_0. I calculate the probability of a case, $J(k^*)$, and the probability of *ex post* compliance, $h(c^*)$. To calculate c^*, I simulate the counterfactual model. I assume that the functional form of J is the exponential distribution and that the functional for of h is the logistic function. I program the numerical simulation such that I can vary the rate parameter of the exponential distribution and the shape and location parameters of the logistic function. To calculate the utility of the government conditional on b_t and b_0, I simulate the economy conditional on those values.

I check convergence to ensure that the optimization routine finds a maximum. The expected utility function is globally concave, which guarantees a unique solution. The solution is interior if $b_0^* > 1$. Finally, I solve the economy conditional on the numerical value of the expected *ex post* trade barriers, $E[b_1^*]$. This produces numerical values for all endogenous parameters in the model.

```
# solve the full model using the analytical solution
solve.full.analytical <- function (var) {

  # utility function
  u.G <- function (b, var) {

    # solve the economy conditional on trade barriers
    var$b <- b
    var <- solve.economy.analytical(var)

```

```
  # calculate utility
  u.G <- - var$w * var$phi.d.star - (1 - var$w) * var$phi.x.star

  # return utility
  return (u.G)
}

# expected utility function
E.u.G <- function(b, var) {

  # solve the compliance stage
  var$b0.star <- b
  var <- solve.compliance(var)

  # calculate expected utility
  E.u.G <- var$pr.case * (var$pr.ex.post.compliance * u.G(var$b.t,
      var) + (1 - var$pr.ex.post.compliance) * u.G(var$b0.star,
     var)) + (1 - var$pr.case) * u.G(var$b0.star, var)

  # return expected utility
  return(E.u.G)
}

# solve the compliance stage
solve.compliance <- function(var) {

  # solve counterfactual model
  var.counterfactual <- solve.counterfactual.analytical(var)

  # equilibrium cost of compliance
  var$c.star <- abs(u.G(var.counterfactual$b.tilde.star, var.
     counterfactual) - u.G(var$b.t, var))

```

```
  # probability of ex post compliance
  var$pr.ex.post.compliance <- 1 / (1 + exp(var$steepness * (var$c.star + var$midpoint)))

  # cut-point for the plaintiff
  var$k.star <- var$pr.ex.post.compliance * (var$b.t - var$b0.star) ^ 2

  # probability of a case
  var$pr.case <- 1 - exp(-var$rate * var$k.star)

  # return solution
  return(var)
}

# solve for the optimal trade barriers
solve.b0.star <- function(var) {

  # optimization
  out <- optim(fn = E.u.G, var = var, par = 2, method = "Brent", lower = 1, upper = 10, control = list(fnscale = -1))
  var$b0.star <- out$par

  # return solution
  return(var)
}

# solve for the optimal ex ante trade barriers
var <- solve.b0.star(var)

# solve the compliance stage in expectation
var <- solve.compliance(var)

```

```
  # expected ex post trade barriers
  var$E.b1.star <- var$pr.case * (var$pr.ex.post.compliance * var$b.
     t + (1 - var$pr.ex.post.compliance) * var$b0.star) + (1 - var$
     pr.case) * var$b0.star

  # probability of enforced compliance
  var$pr.enforced.compliance <- var$pr.case * var$pr.ex.post.
     compliance

  # solve the economy conditional on expected ex post trade barriers
  var$b <- var$E.b1.star
  var <- solve.economy.analytical(var)
  var$b <- NULL

  # return solution
  return(var)
}
```

B.5 Computational Proof of Result 1

I demonstrate Result 1 using a Monte Carlo simulation. The model is nonlinear, so the derivative of a given endogenous variable (e.g., φ_d^* and φ_x^*) with respect to a given exogenous parameter (e.g., b or θ) depends on the values of all other exogenous parameters. I use a Monte Carlo simulation to demonstrate that the derivatives of interest are strictly monotonic and consistent in sign throughout the parameter space.

I conduct one simulation for each exogenous parameter of interest. For each simulation, I run 10,000 iterations. Thus, I calculate each derivative of interest conditional on 10,000 independent random draws from the parameter space. Given the high dimensionality of the parameter space, drawing random samples from the parameter space is far more computationally efficient than a grid search of the parameter space.

I start by defining the sample space. For each exogenous parameter x, I define a uniform probability distribution, $x \sim \mathcal{U}(\underline{x}, \overline{x})$, and choose boundary values, $\underline{x}$ and $\overline{x}$. At each iteration of the simulation, I draw a new random value for each exogenous parameter in the model. The condition $\theta > \sigma - 1$ is required for the average productivity of firms in equilibrium to be finite. If this condition is not satisfied, the integral that yields average productivity is divergent and an equilibrium does not exist. Thus, I set θ equal to $\sigma - 1 + \epsilon$, where $\epsilon \sim \mathcal{U}(0, \overline{\epsilon})$ is a non-negative random shock.

At each iteration, I draw random values for each exogenous parameter, solve the economy numerically, and calculate the numerical derivative of each endogenous parameter of interest y with respect to the exogenous parameter of interest x. The numerical derivative of y with respect to x is $\frac{\partial y^*(x)}{\partial x} \approx \frac{y^*(x+\epsilon) - y^*(x-\epsilon)}{2\epsilon}$, where ϵ is a constant that is small relative to the rate of change. I only calculate derivatives for combinations of parameter values that yield interior solutions.

To reduce computation time, I use the analytical solution for the economy. I also compute the iterations in parallel across multiple cores.

```
# function to solve the economy using the analytical solution
solve.economy.analytical <- function(var) {

  # solve for the domestic production cut-point
  var$phi.d.star <- (var$sigma * var$f * var$M * var$theta  * (1 +
     var$n * var$b ^ (-var$theta)) / (var$I * (1 - var$sigma + var$
     theta))) ^ (1 / var$theta)

  # identify boundary solutions
  var$phi.d.star <- max(var$phi.d.star, 1)

  # solve for the exporting cut-point
  var$phi.x.star <- var$b * var$phi.d.star

```

```
  # solve for the mass of firms that produce for the domestic market
  var$M.d.star <- var$phi.d.star ^ (-var$theta) * var$M

  # solve for the mass of firms that produce for foreign markets
  var$M.x.star <- var$phi.x.star ^ (-var$theta) * var$M

  # solve for the price index
  var$P.star <- (var$sigma / (var$sigma - 1)) * (var$sigma * var$f /
      var$I) ^ ((var$sigma - var$theta - 1) / (var$theta * (1 - var$
     sigma))) * ((var$M * var$theta * (1 + var$n * var$b ^ (-var$
     theta))) / (1 - var$sigma + var$theta)) ^ (-1 / var$theta)

  # solve for consumer welfare
  var$W.star <- 1 / var$P.star

  # return solution
  return(var)
}

# function to sample the equilibrium space
draw.sample <- function(draws) {

  # make a data frame
  sample <- data.frame(iteration = 1:draws)

  # draw parameter values
  sample$n <- 3
  sample$sigma = runif(draws, 1.1, 5)
  sample$theta = sample$sigma + 0.05 + runif(draws, 0, 5) - 1
  sample$f = runif(draws, 1, 10)
  sample$b = runif(draws, 1, 10)
  sample$M = runif(draws, 1, 1000)
  sample$I = runif(draws, 1, 1000)
```

```
  # return the data frame
  return(sample)
}

# function to initialize model
initialize.parameters <- function(x) {

  # convert vector to a named list
  x <- list(n = as.numeric(x[2]),
            sigma = as.numeric(x[3]),
            theta = as.numeric(x[4]),
            f = as.numeric(x[5]),
            b = as.numeric(x[6]),
            M = as.numeric(x[7]),
            I = as.numeric(x[8]))

  # return list
  return(x)
}

# convert a data frame row to a numeric vector with labels
row.to.vector <- function(x) {

  # store parameter names
  names <- names(x)

  # convert data frame to numeric vector
  x <- as.numeric(x)

  # assign parameter names to vector
  names(x) <- names

```

```
  # return
  return(x)
}

# convert list of endogenous parameter values to a named vector
structure.output <- function(x) {

  # assign list elements to a vector
  x <- c(x$phi.d.star,
           x$phi.x.star,
           x$P.star,
           x$W.star)

  # name elements
  names(x) <- c("phi.d.star", "phi.x.star", "P.star", "W.star")

  # return a vector
  return(x)
}

# function to calculate one iteration of the model
run.iteration <- function(exog, param, epsilon) {

  # convert data frame row to a vector
  exog <- row.to.vector(exog)

  # get starting values
  var <- initialize.parameters(exog)
  var.l <- initialize.parameters(exog)
  var.h <- initialize.parameters(exog)

  # exogenous parameter of interest
  var.l[[param]] <- var.l[[param]] - epsilon
```

```
  var.h[[param]] <- var.h[[param]] + epsilon

  # solve economy
  endog <- solve.economy.analytical(var)
  endog.l <- solve.economy.analytical(var.l)
  endog.h <- solve.economy.analytical(var.h)

  # structure output
  endog <- structure.output(endog)
  endog.l <- structure.output(endog.l)
  endog.h <- structure.output(endog.h)

  # calculate comparative statics
  cs <- sign((endog.h - endog.l) / (2 * epsilon))

  # label parameters
  names(cs) <- str_c("delta.", names(cs), sep = "")

  # combine output
  out <- c(exog, endog, cs)

  # return output
  return(out)
}

# function to run the simulation in parallel
run.simulation <- function(sample, param, epsilon) {

  # number of iterations
  iterations <- nrow(sample)

  # initialize progress bar
  pb <- txtProgressBar(max = iterations, style = 3)
```

```
  progress <- function(n) {
    setTxtProgressBar(pb, n)
  }
  opts <- list(progress = progress)

  # run loop in parallel
  out <- foreach(i = 1:iterations,
                 .combine = "rbind",
                 .errorhandling = "stop",
                 .packages = c("stringr"),
                 .options.snow = opts,
                 .export = c("run.iteration", "initialize.parameters
                     ", "solve.economy.analytical", "structure.output
                     ", "row.to.vector")) %dopar%
                 {
                   run.iteration(sample[i,], param = param, epsilon
                       = epsilon)
                 }

  # close the progress bar
  close(pb)

  # convert to data frame
  out <- as.data.frame(out)

  # identify boundary solutions
  out$boundary <- as.numeric(out$phi.d.star == 1)

  # return output
  return(out)
}

# function to summarize comparative statics
```

```
summary <- function(x) {

  # calculate column means
  x <- colMeans(x[,str_detect(names(x), "^delta")])

  # return means
  return(x)
}

# register clusters
cl <- makeCluster(4)
registerDoSNOW(cl)

# sample parameter space
sample <- draw.sample(draws = 10000)

# numerically calculate comparative statics
sim <- run.simulation(sample = sample, param = "b", epsilon = 1e-2)

# keep interior solutions
sim <- filter(sim, boundary == 0)

# comparative statics
summary(sim)

# sample parameter space
sample <- draw.sample(draws = 10000)

# numerically calculate comparative statics
sim <- run.simulation(sample = sample, param = "theta", epsilon = 1e
    -2)

# keep interior solutions
```

```
sim <- filter(sim, boundary == 0)

# comparative statics
summary(sim)
```

B.6 Computational Proof of Result 2

I demonstrate Result 2 using a Monte Carlo simulation of the counterfactual model. In some regions of the parameter space, small changes in the numerical values of exogenous parameters can push the numerical values of some endogenous parameters asymptotically against a lower boundary of 0. When parameters are sufficiently close to 0, numerical precision problems make it difficult to accurately determine the signs of comparative statics. To avoid this problem, I exclude iterations in which one or more parameters is smaller than 1×10^{-5}. As before, I also exclude iterations that involve boundary solutions.

```
# function to solve the economy using the analytical solution
solve.economy.analytical <- function(var) {

  # solve for the domestic production cut-point
  var$phi.d.star <- (var$sigma * var$f * var$M * var$theta  * (1 + var$n * var$b ^ (-var$theta)) / (var$I * (1 - var$sigma + var$theta))) ^ (1 / var$theta)

  # identify boundary solutions
  var$phi.d.star <- max(var$phi.d.star, 1)

  # solve for the exporting cut-point
  var$phi.x.star <- var$b * var$phi.d.star

  # solve for the mass of firms that produce for the domestic market
  var$M.d.star <- var$phi.d.star ^ (-var$theta) * var$M
```

```

  # solve for the mass of firms that produce for foreign markets
  var$M.x.star <- var$phi.x.star ^ (-var$theta) * var$M

  # solve for the price index
  var$P.star <- (var$sigma / (var$sigma - 1)) * (var$sigma * var$f /
       var$I) ^ ((var$sigma - var$theta - 1) / (var$theta * (1 - var$
     sigma))) * ((var$M * var$theta * (1 + var$n * var$b ^ (-var$
     theta))) / (1 - var$sigma + var$theta)) ^ (-1 / var$theta)

  # solve for consumer welfare
  var$W.star <- 1 / var$P.star

  # return solution
  return(var)
}

# function to solve the counterfactual model using the analytical
    solution
solve.counterfactual.analytical <- function(var) {

  # calculate optimal trade barriers
  var$b.tilde.star <- (-(var$w - 1) / (var$n * var$w)) ^ (1 / (-1 -
     var$theta))
  var$b <- var$b.tilde.star

  # solve the economy
  var <- solve.economy.analytical(var)
  var$b <- NULL

  # return solution
  return(var)
}
```

```
# function to sample the equilibrium space
draw.sample <- function(draws) {

  # make a data frame
  sample <- data.frame(iteration = 1:draws)

  # draw parameter values
  sample$n <- 3
  sample$sigma = runif(draws, 1.1, 5)
  sample$theta = sample$sigma + 0.05 + runif(draws, 0, 5) - 1
  sample$f = runif(draws, 1, 10)
  sample$b = runif(draws, 1, 10)
  sample$M = runif(draws, 1, 1000)
  sample$I = runif(draws, 1, 1000)
  sample$w = runif(draws, 0, 1)

  # return a data frame
  return(sample)
}

# function to initialize model
initialize.parameters <- function(x) {

  # convert vector to a named list
  x <- list(n = as.numeric(x[2]),
            sigma = as.numeric(x[3]),
            theta = as.numeric(x[4]),
            f = as.numeric(x[5]),
            b = as.numeric(x[6]),
            M = as.numeric(x[7]),
            I = as.numeric(x[8]),
            w = as.numeric(x[9]))
```

```

  # return a list
  return(x)
}

# convert a data frame row to a numeric vector with labels
row.to.vector <- function(x) {

  # store parameter names
  names <- names(x)

  # convert data frame to numeric vector
  x <- as.numeric(x)

  # assign parameter names to vector
  names(x) <- names

  # return
  return(x)
}

# convert list of endogenous parameter values to a named vector
structure.output <- function(x) {

  # assign list elements to a vector
  x <- c(x$phi.d.star,
         x$phi.x.star,
         x$P.star,
         x$W.star,
         x$b.tilde.star)

  # name elements
```

```
  names(x) <- c("phi.d.star", "phi.x.star", "P.star", "W.star", "b.
      tilde.star")

  # return a vector
  return(x)
}

# function to calculate one iteration of the model
run.iteration <- function(exog, param, epsilon) {

  # convert data frame row to a vector
  exog <- row.to.vector(exog)

  # get starting values
  var <- initialize.parameters(exog)
  var.l <- initialize.parameters(exog)
  var.h <- initialize.parameters(exog)

  # exogenous parameter of interest
  var.l[[param]] <- var.l[[param]] - epsilon
  var.h[[param]] <- var.h[[param]] + epsilon

  # solve economy
  endog <- solve.counterfactual.analytical(var)
  endog.l <- solve.counterfactual.analytical(var.l)
  endog.h <- solve.counterfactual.analytical(var.h)

  # structure output
  endog <- structure.output(endog)
  endog.l <- structure.output(endog.l)
  endog.h <- structure.output(endog.h)

  # calculate comparative statics
```

```
  cs <- sign((endog.h - endog.l) / (2 * epsilon))

  # label parameters
  names(cs) <- str_c("delta.", names(cs), sep = "")

  # combine output
  out <- c(exog, endog, cs)

  # return
  return(out)
}

# function to run the simulation in parallel
run.simulation <- function(sample, param, epsilon) {

  # number of iterations
  iterations <- nrow(sample)

  # initialize progress bar
  pb <- txtProgressBar(max = iterations, style = 3)
  progress <- function(n) {
    setTxtProgressBar(pb, n)
  }
  opts <- list(progress = progress)

  # run loop in parallel
  out <- foreach(i = 1:iterations,
                 .combine = "rbind",
                 .errorhandling = "stop",
                 .packages = c("stringr"),
                 .options.snow = opts,
                 .export = c("run.iteration", "initialize.parameters
                    ", "solve.economy.analytical", "solve.
```

```
                    counterfactual.analytical", "structure.output",
                    "row.to.vector")) %dopar%
                  {
                    run.iteration(sample[i,], param = param, epsilon
                      = epsilon)
                  }

  # close the progress bar
  close(pb)

  # convert to data frame
  out <- as.data.frame(out)

  # return
  return(out)
}

# function to check equilibrium conditions
check.conditions <- function(out, epsilon) {

  # identify boundary solutions
  out$boundary <- out$phi.d.star < 1 + epsilon | out$b.tilde.star <
     1 + epsilon

  # identify precision problems
  out$precision <- out$phi.d.star < epsilon |
    out$phi.x.star < epsilon |
    out$P.star < epsilon |
    out$W.star < epsilon

  # return
  return(out)
}
```

```
# function to summarize comparative statics
summary <- function(x) {

  # calculate column means
  x <- colMeans(x[,str_detect(names(x), "^delta")])

  # return
  return(x)
}

# register clusters
cl <- makeCluster(4)
registerDoSNOW(cl)

# sample parameter space
sample <- draw.sample(draws = 10000)

# numerically calculate comparative statics
sim <- run.simulation(sample = sample, param = "theta", epsilon = 1e
    -2)

# check conditions
sim <- check.conditions(sim, epsilon = 1e-5)
sim <- filter(sim, !boundary & !precision)

# comparative statics
summary(sim)
```

B.7 Computational Proof of Results 3–6

I demonstrate Results 3–6 using a Monte Carlo simulation of the full model. As before, to avoid numerical precision problems, I exclude iterations in which one or more parameters is smaller than 1×10^{-5} and iterations that involve boundary solutions.

Some of the comparative statics in the full model are non-monotonic. Thus, at each iteration, instead of calculating a numerical derivative, I calculate the model conditional on three values of θ — a low value, a middle value, and a high value. The low value is $\theta_L = \sigma - 1 + \epsilon$, the middle value is $\theta_M = \theta_L + \epsilon$, and the high value is $\theta_H = \theta_M + \epsilon$, where $\epsilon \sim \mathcal{U}(0, \bar{\epsilon})$ is a non-negative random shock. I compare the values of the endogenous parameters of interest conditional on these three values of θ to sign the comparative statics. I program the computational algorithm to identify increasing, decreasing, concave, and convex relationships.

```
# function to solve the economy using the analytical solution
solve.economy.analytical <- function(var) {

  # solve for the domestic production cut-point
  var$phi.d.star <- (var$sigma * var$f * var$M * var$theta  * (1 + var$n * var$b ^ (-var$theta)) / (var$I * (1 - var$sigma + var$theta))) ^ (1 / var$theta)

  # identify boundary solutions
  var$phi.d.star <- max(var$phi.d.star, 1)

  # solve for the exporting cut-point
  var$phi.x.star <- var$b * var$phi.d.star

  # solve for the mass of firms that produce for the domestic market
  var$M.d.star <- var$phi.d.star ^ (-var$theta) * var$M
```

```

  # solve for the mass of firms that produce for foreign markets
  var$M.x.star <- var$phi.x.star ^ (-var$theta) * var$M

  # solve for the price index
  var$P.star <- (var$sigma / (var$sigma - 1)) * (var$sigma * var$f / var$I) ^ ((var$sigma - var$theta - 1) / (var$theta * (1 - var$sigma))) * ((var$M * var$theta * (1 + var$n * var$b ^ (-var$theta))) / (1 - var$sigma + var$theta)) ^ (-1 / var$theta)

  # solve for consumer welfare
  var$W.star <- 1 / var$P.star

  # return solution
  return(var)
}

# function to solve the counterfactual model using the analytical solution
solve.counterfactual.analytical <- function(var) {

  # calculate optimal trade barriers
  var$b.tilde.star <- (-(var$w - 1) / (var$n * var$w)) ^ (1 / (-1 - var$theta))
  var$b <- var$b.tilde.star

  # solve the economy
  var <- solve.economy.analytical(var)
  var$b <- NULL

  # return solution
  return(var)
}
```

```

# solve the full model using the analytical solution
solve.full.analytical <- function (var) {

  # utility function
  u.G <- function (b, var) {

    # solve the economy conditional on trade barriers
    var$b <- b
    var <- solve.economy.analytical(var)

    # calculate utility
    u.G <- - var$w * var$phi.d.star - (1 - var$w) * var$phi.x.star

    # return utility
    return (u.G)
  }

  # expected utility function
  E.u.G <- function(b, var) {

    # solve the compliance stage
    var$b0.star <- b
    var <- solve.compliance(var)

    # calculate expected utility
    E.u.G <- var$pr.case * (var$pr.ex.post.compliance * u.G(var$b.t,
        var) + (1 - var$pr.ex.post.compliance) * u.G(var$b0.star,
        var)) + (1 - var$pr.case) * u.G(var$b0.star, var)

    # return expected utility
    return(E.u.G)
  }
```

```

# solve the compliance stage
solve.compliance <- function(var) {

  # solve counterfactual model
  var.counterfactual <- solve.counterfactual.analytical(var)

  # equilibrium cost of compliance
  var$c.star <- abs(u.G(var.counterfactual$b.tilde.star, var.counterfactual) - u.G(var$b.t, var))

  # probability of ex post compliance
  var$pr.ex.post.compliance <- 1 / (1 + exp(var$steepness * (var$c.star + var$midpoint)))

  # cut-point for the plaintiff
  var$k.star <- var$pr.ex.post.compliance * (var$b.t - var$b0.star) ^ 2

  # probability of a case
  var$pr.case <- 1 - exp(-var$rate * var$k.star)

  # return solution
  return(var)
}

# solve for the optimal trade barriers
solve.b0.star <- function(var) {

  # optimization
  out <- optim(fn = E.u.G, var = var, par = 2, method = "Brent", lower = 1, upper = 10, control = list(fnscale = -1))
  var$b0.star <- out$par
```

```

    # return solution
    return(var)
  }

  # solve for the optimal ex ante trade barriers
  var <- solve.b0.star(var)

  # solve the compliance stage in expectation
  var <- solve.compliance(var)

  # expected ex post trade barriers
  var$E.b1.star <- var$pr.case * (var$pr.ex.post.compliance * var$b.
      t + (1 - var$pr.ex.post.compliance) * var$b0.star) + (1 - var$
      pr.case) * var$b0.star

  # probability of enforced compliance
  var$pr.enforced.compliance <- var$pr.case * var$pr.ex.post.
      compliance

  # solve the economy conditional on expected ex post trade barriers
  var$b <- var$E.b1.star
  var <- solve.economy.analytical(var)
  var$b <- NULL

  # return solution
  return(var)
}

# function to compare the full model to the counterfactual model
model.comparison <- function(var)
{
  # solve counterfactual model
```

```
  var.counterfactual <- solve.counterfactual.analytical(var)

  # calculate effects
  var$effect.b0.star <- var$b0.star - var.counterfactual$b.tilde.
     star
  var$effect.E.b1.star <- var$E.b1.star - var.counterfactual$b.tilde
     .star
  var$effect.phi.d.star <- var$phi.d.star - var.counterfactual$phi.d
     .star
  var$effect.phi.x.star <- var$phi.x.star - var.counterfactual$phi.x
     .star
  var$effect.W.star <- var$W.star - var.counterfactual$W.star

  # return
  return(var)
}

# function to sample the equilibrium space
draw.sample <- function(draws) {

  # make a data frame
  sample <- data.frame(iteration = 1:draws)

  # draw parameter values
  sample$n <- 3
  sample$sigma = runif(draws, 1.1, 5)
  sample$theta = sample$sigma + 0.05 + runif(draws, 0, 5) - 1
  sample$f = runif(draws, 1, 10)
  sample$b = runif(draws, 1, 10)
  sample$M = runif(draws, 1, 1000)
  sample$I = runif(draws, 1, 1000)
  sample$w = runif(draws, 0, 1)
  sample$b.t <- 1
```

```
  sample$rate <- runif(draws, 0, 5)
  sample$steepness <- runif(draws, 0, 3)
  sample$midpoint <- runif(draws, -0.25, 0.25)

  # return a data frame
  return(sample)
}

# function to initialize model
initialize.parameters <- function(x) {

  # convert vector to a named list
  x <- list(n = as.numeric(x[2]),
            sigma = as.numeric(x[3]),
            theta = as.numeric(x[4]),
            f = as.numeric(x[5]),
            b = as.numeric(x[6]),
            M = as.numeric(x[7]),
            I = as.numeric(x[8]),
            w = as.numeric(x[9]),
            b.t = as.numeric(x[10]),
            rate = as.numeric(x[11]),
            steepness = as.numeric(x[12]),
            midpoint = as.numeric(x[13]))

  # return a list
  return(x)
}

# convert a data frame row to a numeric vector with labels
row.to.vector <- function(x) {

  # store parameter names
```

```
  names <- names(x)

  # convert data frame to numeric vector
  x <- as.numeric(x)

  # assign parameter names to vector
  names(x) <- names

  # return
  return(x)
}

# convert list of endogenous parameter values to a named vector
structure.output <- function(x) {

  # assign list elements to a vector
  x <- c(x$phi.d.star,
         x$phi.x.star,
         x$P.star,
         x$W.star,
         x$c.star,
         x$b0.star,
         x$E.b1.star,
         x$pr.ex.post.compliance,
         x$pr.case,
         x$effect.b0.star,
         x$effect.E.b1.star,
         x$effect.phi.d.star,
         x$effect.phi.x.star,
         x$effect.W.star)

  # name elements
  names(x) <- c("phi.d.star", "phi.x.star", "P.star", "W.star",
```

```
                "c.star", "b0.star", "E.b1.star", "pr.ex.post.compliance", "pr.case",
                "effect.b0.star", "effect.E.b1.star", "effect.phi.d.star", "effect.phi.x.star", "effect.W.star")

  # return a vector
  return(x)
}

# function to calculate one iteration of the model
run.iteration <- function(exog) {

  # convert data frame row to a vector
  exog <- row.to.vector(exog)

  # get starting values
  var.l <- initialize.parameters(exog)
  var.m <- initialize.parameters(exog)
  var.h <- initialize.parameters(exog)

  # exogenous parameter of interest
  var.l$theta <- var.l$sigma - 1 + runif(1, 0, 3)
  var.m$theta <- var.l$theta + runif(1, 0, 3)
  var.h$theta <- var.m$theta + runif(1, 0, 3)

  # solve economy
  endog.l <- solve.full.analytical(var.l)
  endog.m <- solve.full.analytical(var.m)
  endog.h <- solve.full.analytical(var.h)

  # model comparisons
  endog.l <- model.comparison(endog.l)
  endog.m <- model.comparison(endog.m)
```

```r
endog.h <- model.comparison(endog.h)

# structure output
endog.l <- structure.output(endog.l)
endog.m <- structure.output(endog.m)
endog.h <- structure.output(endog.h)

# create an empty vector
cs <- rep("missing", length(endog.l))
names(cs) <- names(endog.l)

# calculate comparative statics
cs[endog.l < endog.m & endog.m > endog.h] <- "concave"
cs[endog.l > endog.m & endog.m < endog.h] <- "convex"
cs[endog.l < endog.m & endog.m < endog.h] <- "increasing"
cs[endog.l > endog.m & endog.m > endog.h] <- "decreasing"

# label parameters
names(cs) <- str_c("delta.", names(cs), sep = "")
names(endog.l) <- str_c("L.", names(endog.l), sep = "")
names(endog.m) <- str_c("M.", names(endog.m), sep = "")
names(endog.h) <- str_c("H.", names(endog.h), sep = "")

# format
exog <- as.data.frame(t(exog), stringsAsFactors = FALSE)
cs <- as.data.frame(t(cs), stringsAsFactors = FALSE)
endog.l <- as.data.frame(t(endog.l), stringsAsFactors = FALSE)
endog.m <- as.data.frame(t(endog.m), stringsAsFactors = FALSE)
endog.h <- as.data.frame(t(endog.h), stringsAsFactors = FALSE)

# combine
out <- cbind(exog, endog.l, endog.m, endog.h, cs)

```

```
# order parameters
out <- select(out,
              n, sigma, theta, f, M, I, w, b.t, rate, steepness, midpoint,
              L.phi.d.star, M.phi.d.star, H.phi.d.star,
              L.phi.x.star, M.phi.x.star, H.phi.x.star,
              L.P.star, M.P.star, H.P.star,
              L.W.star, M.W.star, H.W.star,
              L.c.star, M.c.star, H.c.star,
              L.b0.star, M.b0.star, H.b0.star,
              L.E.b1.star, M.E.b1.star, H.E.b1.star,
              L.pr.ex.post.compliance, M.pr.ex.post.compliance, H.pr.ex.post.compliance,
              L.pr.case, M.pr.case, H.pr.case,
              L.effect.b0.star, M.effect.b0.star, H.effect.b0.star,
              L.effect.E.b1.star, M.effect.E.b1.star, H.effect.E.b1.star,
              L.effect.phi.d.star, M.effect.phi.d.star, H.effect.phi.d.star,
              L.effect.phi.x.star, M.effect.phi.x.star, H.effect.phi.x.star,
              L.effect.W.star, M.effect.W.star, H.effect.W.star,
              delta.c.star, delta.b0.star, delta.E.b1.star,
              delta.pr.ex.post.compliance, delta.pr.case,
              delta.phi.d.star, delta.phi.x.star, delta.P.star, delta.W.star,
              delta.effect.b0.star, delta.effect.E.b1.star,
              delta.effect.phi.d.star, delta.effect.phi.x.star, delta.effect.W.star)

# return
return(out)
```

```
}

# function to run the simulation in parallel
run.simulation <- function(sample) {

  # number of iterations
  iterations <- nrow(sample)

  # initialize progress bar
  pb <- txtProgressBar(max = iterations, style = 3)
  progress <- function(n) {
    setTxtProgressBar(pb, n)
  }
  opts <- list(progress = progress)

  # run loop in parallel
  out <- foreach(i = 1:iterations,
                 .combine = "rbind",
                 .errorhandling = "stop",
                 .packages = c("stringr", "dplyr"),
                 .options.snow = opts,
                 .export = c("run.iteration", "initialize.parameters
                    ", "solve.economy.analytical", "solve.
                    counterfactual.analytical", "solve.full.
                    analytical", "model.comparison", "structure.
                    output", "row.to.vector")) %dopar%
                 {
                   run.iteration(sample[i,])
                 }

  # close the progress bar
  close(pb)

```

```
  # convert to data frame
  out <- as.data.frame(out)

  # return
  return(out)
}

# function to check equilibrium conditions
check.conditions <- function(x, epsilon) {

  # extract numeric columns
  temp <- select_if(x, is.numeric)

  # loop through iterations
  for(i in 1:nrow(temp)) {

    # mark if any parameters are sufficiently small
    x$drop[i] <- sum(abs(temp[i,]) < epsilon) > 0
  }

  # drop observations
  x <- filter(x, !drop)

  # return
  return(x)
}

###############################################
# calculate comparative statics
###############################################

# register clusters
cl <- makeCluster(8)
```

```
registerDoSNOW(cl)

# sample parameter space
sample <- draw.sample(draws = 10000)

# numerically calculate comparative statics
sim <- run.simulation(sample = sample)

# check conditions
sim <- check.conditions(sim, epsilon = 1e-5)

# result 3, figure 5
table(sign(sim$M.effect.b0.star)) # decreasing
table(sign(sim$M.effect.E.b1.star)) # decreasing

# result 4, figure 6
table(sim$delta.b0.star) # decreasing
table(sim$delta.E.b1.star) # decreasing
table(sim$delta.effect.b0.star) # convex
table(sim$delta.effect.E.b1.star) # convex

# figure 7
table(sim$delta.c.star) # decreasing
table(sim$delta.pr.ex.post.compliance) # increasing
table(sim$delta.pr.case) # concave

# result 5, figure 8
table(sign(sim$M.effect.phi.d.star)) # increasing
table(sign(sim$M.effect.phi.x.star)) # decreasing

# result 6, figure 9
table(sim$delta.phi.d.star) # decreasing
table(sim$delta.phi.x.star) # decreasing
```

```
table(sim$delta.W.star) # decreasing
table(sim$delta.effect.phi.d.star) # concave
table(sim$delta.effect.phi.x.star) # convex
table(sim$delta.effect.W.star) # concave
```