The Welfare Benefits of Raising your Standards: Evidence from International Regulatory Protectionism

Luca Macedoni* Ariel Weinberger†
Aarhus University University of Oklahoma

February 6, 2018

VERY PRELIMINARY AND INCOMPLETE – PLEASE DO NOT CITE

Abstract

We study the effects of non-tariff technical measures, in the form of minimum quality standards, on the welfare of consumers. Technical measures have become an increasingly relevant trade policy tool and a key part of modern trade agreements. While the literature considers these measures equivalent to trade costs, in this paper, we build a model of heterogeneous, monopolistically competitive firms to identify the circumstances under which minimum quality standards can be welfare improving. Raising a country's quality standard improves welfare by reducing the excessive entry of low-quality firms, which are otherwise able to serve an economy by lowering their markups. Trade openness reduces the need for a quality standard: in a two-country framework, the lower the trade costs are, the lower the optimal minimum quality standard. The predictions of the model are confirmed in several ways. We examine entry and exit dynamics of European firms by size class, and then turn to sales distributions using only Chilean firms. In both cases we test the response to the prevalence of industry-specific regulatory protectionism. Lastly, we provide a quantitative analysis of optimal standards and compare these to the measures of standards in the data.

JEL Classification: F12, F13.

Keywords: Allocative Efficiency, Minimum Quality Standards, Variable Markups, Non-Homothetic Preferences.

*Address: Department of Economics and Business Economics, Aarhus University, Fuglesangs Allé 4, 8210 Aarhus V. E-mail: lmacedoni@econ.au.dk

†Address: Department of Economics, University of Oklahoma. 338 Cate Center Drive, Norman, OK 73019. E-mail: Ariel.Weinberger@ou.edu

1 Introduction

The secular decline in tariffs and in traditional trade policy measures such as quotas has made non-tariff technical measures increasingly relevant and subject of public debates (Maskus et al., 2000; Baldwin et al., 2000). Technical measures, or regulatory protectionism (Baldwin et al., 2000), consist of sanitary and phythosanitary standards, technical barriers to trade, and pre-shipment inspections and other formalities (UNCTAD, 2012). The literature has traditionally considered technical measures as a form of trade costs (Baldwin et al., 2000; Chen and Novy, 2011; Fontagné et al., 2015) – ones that primarily impact the extensive margin – thus implicitly equating technical measures to protectionism. In this study, we explore an angle that is new to this literature, namely whether allocatively inefficient markets provide a rationale for quality standards. If the market allocation results in an inefficient number of producers, regulations can be applied in a way that reduces distortions. We provide an answer to two crucial questions: 1) can technical measures improve welfare? and 2) does trade openness affect the optimal degree of regulatory protectionism?

To answer to these questions, we introduce, to a standard model of trade with heterogeneous firms (Melitz, 2003), a particular form of technical measure: minimum quality standards. Minimum quality standards are a vertical norm (Baldwin et al., 2000), meaning that they can be easily characterized as being more or less stringent, such as restrictions on cars' emissions or food safety[1]. Firms are monopolistically competitive and differ in terms of the exogenously determined quality of their products (Feenstra and Romalis, 2014). The higher the quality of a firm, the higher its sales and markup. The market determines a minimum level of quality: consumers would not purchase goods below a specific cutoff. We model the minimum quality standard as an exogenous cutoff, larger than the market determined one. The standard forces low quality firms out of the market, which is consistent with the empirical evidence[2].

The model features consumers with non-homothetic preferences of the Indirect Additive (IA) class (Bertoletti et al., 2017). Such preferences allow for a tractable analysis of the welfare effects of quality standards and provide a parsimonious way to build a model consistent with several features of sales, productivity and markup distributions of the data. Analogously to other models that feature non-homothetic preferences, our model predicts that a market allocation would be inefficient. The inefficiency takes the form of a *business*

[1] Horizontal norms arise when the local firms' differentiated good is adopted as a norm, as electric plugs (Baldwin et al., 2000).

[2] Technical measures to trade mainly affect the extensive margin of exports, rather than the intensive margin (Fontagné et al., 2015; Fernandes et al., 2015; Ferro et al., 2015). Therefore, we abstract from variable costs of quality control or inspection associated with the quality standard or quality upgrading.

stealing bias: too many low-quality firms are active as they do not internalize the downward pressure imposed by their production on the prices of other firms. Consumers are willing to purchase low-quality goods provided that markups are low enough in the laissez faire economy. The business stealing bias dominates another distortion commonly labeled "lack of appropriability", which generates *too little* production from low-quality firms, and occurs when firms cannot fully seize or appropriate the gains from a new variety[3].

Generally, the allocation decision involves two choices: the quantity produced by each firm, and the number of firms that attempt to enter the market. We are able to focus our attention on the allocation of market share among producing firms by making an assumption common to the literature with firm heterogeneity, that firms draw their quality from a Pareto distribution (Chaney, 2008; Arkolakis et al., 2012, 2015). The Pareto assumption, combined with the absence of fixed production costs, constrains entry to be proportional to population and the fixed entry cost. While the Pareto assumption eliminates inefficiency from entry, it allows for a tractable examination of the distortion present in the production allocation[4]. This type of misallocation has been highlighted by the aggregate productivity literature going back to Basu and Fernald (2002), with more recent quantitative evidence in Restuccia and Rogerson (2008) and Hsieh and Klenow (2009). In the trade literature, distortions due to markup variability have also been shown to be quantitatively significant (Edmond et al., 2015; Weinberger, 2017).

The first contribution of this paper is to show that a minimum quality standard improves the efficiency of the market allocation. In a closed economy, a minimum quality standard has three effects on the welfare of consumers. First, by forcing the products of low-quality firms out of the market, it raises the average quality in the economy, thus, improving welfare. However, the second effect of the quality standard is a reduction of the number of varieties available for consumption. In models featuring love for variety (Krugman, 1980), as ours, fewer varieties reduce welfare. Third, the minimum quality standard effectively reduces the business stealing bias, bringing the allocation of the market economy closer to the efficient allocation. The model generates a non-monotone, hump-shaped relationship between the quality standard and welfare. For low levels of the quality standard, reducing the business stealing bias and improving the average quality dominate the welfare loss from diminishing variety. If the standard becomes too restrictive — above its optimal level — the welfare loss from diminishing variety would offset the welfare enhancing components of the standard.

The welfare improving property of quality standards hinges on the presence of a distor-

---

[3]This intuition is present in Mankiw and Whinston (1986) and Dhingra and Morrow (2016)

[4]Because of business stealing, the market-determined quality cutoff is too low relative to the efficient allocation: there is an excessive number of active low-quality firms.

tion, the business stealing bias, whereby too many low-quality firms are active in a market. In models of monopolistically competitive, heterogeneous firms, the presence of such a distortion depends on the presence of variable markups, which in turn are generated by the properties of demand. Under traditional CES preferences and constant markups the market allocation is optimal because, for each variety, the business stealing bias is equal to the gain in utility from an increase in variety (Feenstra and Kee, 2008; Dhingra and Morrow, 2016)[5]. We extend our framework to several examples with variable markups to provide a more general understanding of when quality standards can raise welfare. Under non-homothetic separable preferences — in an extension, we take the Stone-Geary type used in Simonovska (2015) as an example — a minimum quality standard raises welfare, but we show that the optimal standard is lower than the baseline IA case. With IA preferences, markups depend *only on income* (Bertoletti and Etro, 2017), while with non-homothetic separable preferences, markups depend on both income and the number of competitors. With separable preferences, the quality standard, by lowering the number of firms, simultaneously reduces competition for the surviving firms. Finally, variable markups are also a feature of the homothetic "Quadratic Mean of Order R" (QMOR) model detailed in Feenstra (2014). In the QMOR, markups vary depending *only on the number of competitors*. In this case, we show that a quality standard reduces welfare.

The second contribution of the paper is to study whether trade openness affects the optimal degree of regulatory protectionism. Consider the following thought experiment: suppose that two symmetric countries have to choose a common quality standard to maximize their welfare. Such a scenario could represent the decision phase in trade agreements or within the European Union. How does the optimal quality standard vary with the variable trade costs associated with exporting? A reduction in trade costs reduces the inefficient allocation between exporters and domestic producers. The mechanism differs from that explored in our first contribution. Instead of removing the low-quality firms, a trade liberalization reallocates market share towards high-quality goods even with the same number of domestic varieties available.[6]

Since the quality standard improves upon an inefficient allocation, lower trade costs reduce the optimal level of the quality standard. Although a low level of the quality standard can be welfare improving in a closed economy, as economies open to trade the optimal

[5] However, Benassy (1996), building on the working paper version of Dixit-Stiglitz, shows that in the presence of a "variety externality", the market allocation would not be efficient under CES preferences. In fact, we show that if there is a negative externality — consumers prefer fewer varieties than the outcome of a market allocation — a minimum quality standard can raise welfare.

[6] We focus only on the quality standard, so that the optimal tariff results of Felbermayr et al. (2013) and Demidova and Rodriguez-Clare (2009) are beyond the scope of this paper.

level decreases. For this reason we emphasize that the findings of this paper are not a call for protectionism, but instead offer support of a dual approach for policymakers: pushing towards lower trade costs while lowering the restrictiveness of quality standards.

We then turn to the data to examine the prevalence of non-tariff measures (NTMs) as a policy counterpart to our theoretical quality standard. NTMs are an important way through which policy makers can limit local market access through the imposition of regulatory measures. As modern trade agreements stray from a focus on traditional protectionism measures such as tariffs and quotas, clear objectives on regulatory standards has become a main focus. Since NTMs have been considered as a form of trade costs that primarily impact the extensive margin (Ferro et al., 2015; Fontagné et al., 2015), we believe they provide a conforming mapping to quality standards.

We measure the degree to which certain products are covered by regulatory protectionism across countries by using a new database that classifies domestic regulations as types of NTMs at the product level for a wide span of WTO countries. The database is a product of an important initiative by UNCTAD to "increase transparency and understanding about regulations and trade control measures" (UNCTAD, 2017). It has put together an exhausting list of measures currently enforced and registered from legal texts. From the perspective of our study, it is important that these regulations cover *all* products within the imposing country, so that they affect imported and domestic goods alike.

Our theoretical framework can be used to check various predictions that are salient to our welfare results. We use the prevalence of NTMs as variation in several contexts. First, we are interested in the entry and exit dynamics that are implied by the imposition of quality standards. In order to study this we turn to a Business Demographics database of EU countries. This database provides information on the number of active firms, birth rates, and exit rates at the industry level and for various size classes for most EU countries. Although we can measure these only at very aggregate industry levels, since the variation is at both the industry and country level, we can run a fixed effects specification that controls for both industry and country characteristics. Furthermore the industries can be matched to the 2 digit HS classification of NTMs that is provided by CEPII and is constructed from the UNCTAD data. Our main finding is that entry rates and death rates are both higher in more regulated industries, as would be predicted by the model. More importantly, the entry and death rates are driven by *small* firms, which are the ones we would expect to be affected by the imposition of standards (assuming these are the low-quality firms).

Second, we turn to the sales distribution implications of standards. Welfare increases when market share is reallocated from low-quality to high-quality firms, which raises the production of high-quality goods and increases average quality – at least for low enough standards such that the loss of variety does not dominate. We test the distributional implications of regulatory protectionism with a census of Chilean firms that report revenues, inputs, and industry classification. The information on input costs and import/export behavior is used to classify firms as high or low quality. Then, we run a difference-in-difference specification to test whether sales distributions are altered by NTMs. We find that the ratio of average sales of high-quality firms relative to the average sales of low-quality firms is higher in industries with a greater number of NTMs. We check whether this holds when we use time variation in NTM imposition and find that our results are robust. The findings are therefore consistent with NTMs reducing the business-stealing bias by allowing higher quality firms to expand.

Finally, we turn to a quantitative analysis of welfare. The model can be calibrated using moments from the Chilean sales distribution, thus allowing us to compute optimal standards and sales distribution parameters at the industry level. We compare these to the prevalence of NTMs found in the data to assess whether the standards imposed by policy makers are in line with the optimal standards implied by the model. [TO BE COMPLETED].

This paper is organized as follows. Sections 2, 3, and 4 cover the theory: we begin by discussing the baseline model and show how the market allocation is inefficient, relative to what a planner would choose. Then, in section 3, we prove that a quality standard can improve upon the market allocation in a closed economy, and examine how the results change under alternative classes of preferences. In section 4, we study how trade openness affects the optimal level of the quality standard. Next, we turn to the empirical analysis. Section 5 describes the newly constructed NTM database as well as our outcome measures. Section 6 describes our empirical specification and results. We conclude in section 7.

2 Market Allocation and Welfare

In this section, we build a theory on the effects of minimum quality standards on the welfare of consumers. We adopt standard supply side assumptions and leverage specific demand side features that allow us to highlight the inefficient allocation that is the central topic of this paper. The model features two main ingredients: quality-differentiated firms and consumers with non-homothetic preferences. As in Hallak and Sivadasan (2013) and Feenstra and Romalis (2014), quality is a firm-specific demand shifter: when a government chooses a minimum quality standard, all firms with a quality level below such standard are forced out of the market. Consumer preferences take a functional form within the IA class detailed by Bertoletti and Etro (2017) and Bertoletti et al. (2017), in which the elasticity of demand depends only on income and prices of other varieties.

As a reference point, we begin the section by describing the market allocation that arises in a closed economy without a minimum quality requirement. We then compare the market allocation to the efficient allocation chosen by a social planner, finding that the market allows for the survival of too many low-quality firms.

2.1 A Model with Non-Homothetic Preferences

We begin the section by describing the consumers' and firms' problem in a general multi-country setting. We then turn to solving for the general equilibrium aggregates and welfare in a closed economy. The section continues with the analysis of the social planner's problem. For more details, we refer the reader to appendix 8.1.

2.1.1 Consumer's Problem

We adopt the linear case of the IA preferences proposed by Bertoletti et al. (2017). The utility of the representative consumer in country j is given by:

$$U_j = \frac{\left[a \int_{\Omega_j} z(\omega) x^c(\omega) dz - 1\right]^2}{2 \int_{\Omega_j} \left(z(\omega) x^c(\omega)\right)^2 dz} \tag{1}$$

where a is a positive constant, $x^c(\omega)$ is the quantity consumed of variety ω, $z(\omega)$ is a variety specific demand shifter, which we interpret as quality, and Ω_j is the set of varieties available for consumption in country j. In the next section we extend our analysis to various classes of preferences – including both non-homothetic and homothetic preferences – where the market allocation is inefficient due to variable markups.[7]. The consumer's budget constraint is:

$$\int_{\Omega_j} p(\omega) x^c(\omega) dz \leq y_j$$

where $p(\omega)$ is the price of variety ω and y_j is per capita income in country j. The consumer chooses $x^c(\omega)$, $\omega \in \Omega_j$, to maximize its utility subject to the budget constraint.

There are L_j identical consumers in country j. Let $x(\omega) = L_j x^c(\omega)$ be the aggregate demand for the variety ω. Solving the consumer's problem and aggregating across consumers yield the following aggregate inverse demand:

$$p(\omega) = y_j a z(\omega) - \frac{y_j^2 \lambda_j}{L_j} z(\omega)^2 x(\omega) \tag{2}$$

[7] In all these cases, the demand elasticity depends on the number of consumed varieties, and only in the non-homothetic cases do they also depend on income.

where λ_j is the marginal utility of income in j.

2.1.2 Firms' Problem

Firms differ in the quality z of their products. Since z denotes the only characteristic of a firm, for convenience, we refer to variety ω produced by a firm with quality z, as variety z. Marginal costs of production and delivery are identical across firms: $w_i \tau_{ij} c$, where w_i is the wage, τ_{ij} is the iceberg trade cost (with $\tau_{ii} = 1$), and c is a cost parameter[8]. As in the standard Melitz (2003) model, in each country there is a pool of potential entrants. Upon entry, firms pay a fixed cost f_E of entry and discover their quality draw z. z is drawn from a Pareto distribution with shape parameter κ and shift parameter b_i. The CDF of the distribution is $G_i(z) = 1 - \left(\frac{b_i}{z}\right)^{\kappa}$, while the pdf is $g_i(z) = \frac{\kappa b_i^{\kappa}}{z^{\kappa+1}}$. Only a mass J_i of firms pays the fixed cost of entry. Free entry drives expected profits equal to the fixed cost of entry.

Given the quality draw z, a firm maximizes its profits in each destination j by choosing quantity $x(z)$ taking λ_j as given. Profits are given by:

$$\pi_{ij}(z) = y_j a z x - \frac{y_j^2 \lambda_j}{L_j} z^2 x^2 - w_i \tau_{ij} c x \tag{3}$$

Profit maximization yields the following optimal quantity:

$$x_{ij}(z) = \frac{L_j a}{2 y \lambda z^2} \left[z - \frac{w_i \tau_{ij} c}{y_j a} \right] \tag{4}$$

There exists a quality level z_{ij}^*, such that $x(z_{ij}^*) = 0$. z_{ij}^* is the market-determined minimum quality level that allows for a positive demand — the quality cutoff. We use the term "market-determined" to distinguish between z_{ij}^* and the minimum quality requirement that a government introduces in the following sections. Setting (4) equal to zero yields:

$$z_{ij}^* = \frac{w_i \tau_{ij} c}{a y_j} \tag{5}$$

For a quality level below the cutoff $z < z^*$, a firm has zero demand in economy j. Thus, the share of entrants that is able to serve to a market is given by $1 - G_i(z_{ij}^*)$. The larger the costs to reach a destination j, the larger the minimum quality allowed. Moreover, the

[8]Several models, for instance those of Hallak and Sivadasan (2013) and Feenstra and Romalis (2014), feature marginal costs that are an increasing function of quality. Adding such a feature to our model, however, does not add any inefficiency that could be improved by the minimum quality standard.

quality cutoff declines with the per capita income of the destination. Using (5) in (4) yields:

$$x_{ij}(z) = \frac{L_j a}{2y\lambda z^2}\left[z - z_{ij}^*\right] \tag{6}$$

As $x_{ij}(z)$ is increasing in z, active firms with higher quality sell larger quantities of their products. As $x_{ij}(z)$ is declining in z_{ij}^*, firms produce less to economies that are more difficult to reach. Substituting (6) and (5) in the inverse demand function (2) yields the following pricing rule:

$$p_{ij}(z) = \frac{y_j a}{2}\left[z + z_{ij}^*\right] = w_i \tau_{ij} c \underbrace{\frac{z + z_{ij}^*}{2z_{ij}^*}}_{\text{Markup}} \tag{7}$$

Markups are increasing in z: higher quality firms charge higher markups.[9].

2.1.3 Market Allocation in a Closed Economy

Let us consider the equilibrium of a closed economy. In equilibrium, expected profits equal the fixed cost of entry f_E, and aggregate revenues equal market size. Without loss of generality we normalize per capita income $y = w$ to one and we leave the derivation to the appendix. The quality cutoff is a constant:

$$z^{B*} = \frac{c}{a}, \tag{8}$$

where we represent the baseline market allocation with a superscript B. The lowest quality allowed by the market is increasing in the marginal cost of production c and decreasing in the preference parameter a. IA preferences feature a constant quality cutoff and, therefore, markups, only depend on real income. Under additively separable preferences, the cutoff additionally depends on the number of competitors (Simonovska, 2015).

Thanks to the Pareto assumption on the distribution of quality draws, the mass of entrants J is proportional to the size of the economy, and inversely proportional to the fixed cost of entry:

$$J^B = \frac{L}{f_E(\kappa + 1)} \tag{9}$$

[9]This is a point made by Bertoletti and Etro (2017) in their description of the IA model. It receives empirical support from Bastos and Silva (2010), Manova and Zhang (2012), Martin (2012) and Dingel (2015).

Finally, the number of varieties available for consumption N is given by:

$$N^B = J^B \left(\frac{b}{z^{B*}}\right)^{\kappa} = \frac{L}{f_E(\kappa+1)} \left(\frac{b}{z^{B*}}\right)^{\kappa}$$

2.2 Allocation of a Social Planner

Is the market allocation efficient? To answer to this question, we consider the problem of a social planner who maximizes the utility of the representative consumer (1) by choosing the quantity of each variety produced, $x^P(z)$, and the mass of potential entrants J^P. We denote the social planner allocation with the superscript P[10].

The social planner faces a resource constraint for the economy such that the labor supply equals the economy-wide fixed cost of entry bill and the total amount of resources used for production:

$$J^P \left[\int_{z^{*P}}^{\infty} cx^P(z)g(z)dz + f_E\right] = L.$$

Relegating the details to the appendix, we show that the efficient supply of a variety z is:

$$x^P(z) = \frac{aL}{\lambda z^2}\left[z - z^{*P}\right], \tag{10}$$

where λ is still the marginal utility of income in the market allocation.

The quantity chosen by the planner differs from the market determined one (6) along two dimensions[11]. First, keeping z^* constant across the two scenarios, the planner chooses a quantity that is twice the market determined one[12]. Firms that are monopolistically competitive have some market power over the differentiated variety they produce, and exercise their market power by restricting the supply of their good, thus increasing markups. Second, there exists a firm with a quality $\tilde{z}$ such that the quantity chosen by the planner for $\tilde{z}$ is identical to the one chosen by the planner. Firms with $z > \tilde{z}$ produce less than what required by an efficient allocation, and firms with $z < \tilde{z}$ produce more than what required by an efficient allocation. In other words, relative to the efficient allocation chosen by the planner, the market leads to overproduction of low-quality goods and underproduction of high-quality goods[13].

[10] The planner would want to choose the quality cutoff z^*, but we verified that z^* is such that $x(z^*) = 0$.

[11] We set $y = 1$ in (6)

[12] $\frac{x^P(z)}{x^B(z)} = 2\left(\frac{z - z^{*P}}{z - z^{*B}}\right)$

[13] This can be shown using $\frac{\partial \frac{x^P(z)}{x^B(z)}}{\partial z} = 2\frac{z^{*P} - z^{*B}}{(z - z^{*B})^2} > 0$, and that $\tilde{z} = 2z^{*P} - z^{*B}$.

We can now turn to the social planner allocation decision. The quality cutoff is given by:

$$z^{P*} = \frac{c}{a}\frac{\kappa+1}{\kappa}. \tag{11}$$

The planner chooses a quality cutoff that is $1+\frac{1}{\kappa}$ times larger the market determined one — which creates a rationale for a minimum quality standard. The mass of entrants $J^P = \frac{L}{f_E(\kappa+1)}$ is identical to the market-determined one (9). This is, in fact, not surprising since the mass of entrants is fixed by the population and fixed entry costs. Arkolakis et al. (2017) and Feenstra (2014) show that the assumptions of no fixed production costs and an unbounded Pareto distribution imply a "proportionality relation" between expected profits from entering the market and expected revenue. Entry is constrained to be proportional to population given our supply side assumptions regardless of preferences, so that the inefficiency in the baseline model lies solely in the selection of firms through the cutoff z^*.

2.2.1 Discussion

As in Dhingra and Morrow (2016) (DM), there are two possible distortions in the model: i) a selection bias in the number of firms that enter the market (making average quality too high or too low), and ii) a quantity bias in terms of the allocation of production for each firm. The first type of distortion is due to lack of appropriability: in making their entry decision, firms do not take into account the social gains from an increase in variety. With "love of variety", the selection bias tends to reduce the number of firms below the social optimum. Firm heterogeneity in market power generates the second distortion: in making their entry decision, firms do not take into account how their entry reduces production and prices of other firms. This "business stealing" effect (DM and Mankiw and Whinston (1986)) raises entry above the optimum because it allows lower quality firms to steal business from high quality firms — which raise their markup as they move up their demand curve. Although the two distortions push the efficient mass of firms in opposite directions, the fact that there is no entry inefficiency means that which bias dominates is pinned down by the quality cutoff. In the IA model, the quality cutoff is too low and there are too many consumed varieties relative to what the social planner would choose.

The two distortions can be summarized by the elasticity of utility and the inverse demand elasticity. The elasticity of utility is defined as: $\epsilon(x(z)) = \frac{u'(x(z))x(z)}{u(x(z))}$, and it represents the contribution to welfare of a new variety that is not captured by real (Vives, 2001). In the language used by DM, we denote $1/\epsilon(x(z))$ as the "social markup" because it is the markup a social planner would choose for every firm. In contrast, the inverse demand elasticity, $\mu(x(z)) = \frac{u''(x(z))x(z)}{u'(x(z))} = \frac{z-z^*}{z+z^*}$, governs firms' markups in the market equilibrium. As in DM, we take the average of each elasticity:

$$\frac{1}{\bar{\epsilon}(x^P(z))} = \left[\frac{\int_{z^*}^{\infty} u'(x^{opt}(z))x^{opt}(z)}{\int_{z^*}^{\infty} u(x^{opt}(z))}\right]^{-1} = \frac{1}{\kappa} \tag{12}$$

$$\bar{\mu}(x^B(z)) = \frac{\int_{z^*}^{\infty} \mu_{ij} r_{ij}(z)\kappa \frac{b^\kappa}{z^{\kappa+1}} dz}{\int_{z^*}^{\infty} r_{ij}(z)\kappa \frac{b^\kappa}{z^{\kappa+1}} dz} = \frac{1}{\kappa+1}, \tag{13}$$

where the average inverse demand elasticity is weighted by revenues. With *too much entry*, it is intuitive that the average market markup is below the social markup as lower per-firm production implies that firms face more elastic demand and there are too many low-quality, low-markup firms. The next section describes how a minimum quality requirement would eliminate these low quality firms, shifting production to high-quality firms.

3 Quality Standard and Welfare

In this section, we introduce a minimum quality standard in the baseline model. As the market allocation is inefficient, the introduction of a minimum quality standard improves welfare by forcing low-quality firms out of the market. Additionally, we extend the analysis to alternative classes of preferences where variable markups cause distortions and establish in which cases a quality standard can improve welfare.

The government of the closed economy sets a minimum quality requirement $\bar{g} \geq z^*$, such that a firm with quality $z < \bar{g}$ is not allowed to sell in the economy. The quality requirement is a vertical norm (Baldwin et al., 2000): $\bar{g}$ can be easily interpreted as more or less restrictive[14]. Thus, we can think of $\bar{g}$ as the minimum level required for a food to be considered safe, or the maximum level of emissions allowed in cars. Since firms' quality is exogenously determined, the policy only affects the selection of firms into the domestic market (Figure 1). In particular, the larger $\bar{g}$ becomes, the more low-quality, low-revenues firms are forced out of the market. The model is consistent with the evidence documented by Fontagné et al. (2015) and Fernandes et al. (2015), who show that more restrictive technical barriers reduce entry and promote exit, especially for the smallest firms[15].

[14] We assume perfect information, thus abstracting from information frictions that could be improved by a standard.

[15] We abstract from compliance and inspection costs, which would reduce the welfare improving effect of the policy. As long as the increase in compliance costs does not entirely offset the welfare improving effect of a minimum quality requirement, the result of this subsection still holds. Quantifying the effect of a quality requirement on compliance costs would require a much greater level of detail both in the theory and in the data, which goes beyond the purpose of this paper.

Figure 1: Minimum Quality Standard and Selection

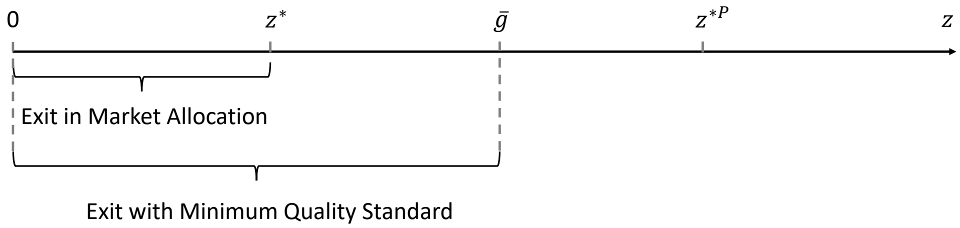

The addilog utility function determines a quality cutoff z^* that is independent of the minimum quality requirement. As in equilibrium, $w = y$, and the market-determined quality cutoff is $z^* = \frac{c}{a}$. The result is in line with the evidence of Ferro et al. (2015), who document no effects of TBT on trade intensity. Moreover, other studies that document a small effect of regulations on export intensity (Fontagné et al., 2015; Fernandes et al., 2015), find the effect on the extensive margin to be quantitatively more relevant.

3.1 Welfare with the Minimum Quality Standard

Let $\bar{z} = \bar{g}$ denote the firm with quality equal to the minimum allowed: only firms with quality $z \in [\bar{z}, \infty)$ are allowed to sell their products in the closed economy. It is convenient to write our variables as a function of $g = \frac{\bar{z}}{z^*} \in [1, \infty)$, a measure of the restrictiveness of the quality standard. If $g \leq 1$, the standard is ineffective: the market-determined quality cutoff z^* is larger or equal to the minimum allowed $\bar{z}$. For $g > 1$, the government is enforcing a higher quality standard than the one determined by the market. The measure g is related to the probability of a firm being active under the restriction, relative to the same probability without the restriction: $\frac{P(z \geq \bar{g} | g > 1)}{P(z \geq \bar{g} | g = 1)} = g^{-\kappa}$.

Leaving the derivations to the appendix, the utility of the representative consumer is a non-monotone function of the quality standard g:

$$U^R = \frac{La^2}{8f_E} \left(\frac{ab}{c} \right)^{\kappa} \frac{\left[g^{-\kappa} + \frac{\kappa}{\kappa+2} g^{-(\kappa+2)} - \frac{2\kappa}{\kappa+1} g^{-(\kappa+1)} \right]^2}{g^{-\kappa} - \frac{\kappa}{\kappa+2} g^{-(\kappa+2)}} \tag{14}$$

The relationship is hump-shaped: in the appendix we prove that the utility is maximized at $g > 1$, and that $g^{opt} = \tilde{\kappa}$, where $\tilde{\kappa}$ is a non-linear function which is declining in κ. The optimal level of the standard $\bar{g}^{opt}$ is then proportional to the market-determined cutoff:

$$\bar{g}^{opt} = g z^* = \tilde{\kappa} \frac{c}{a} \tag{15}$$

Hence, less productive economies (lower c) have higher optimal quality standards, while more productive economies would choose lower levels of $\bar{g}$. Figure 2 shows the relationship between welfare and the minimum quality standard. For $\kappa = 4$, welfare is maximized at $g = 1.32$: the government sets a standard which reduces the probability of a firm being active by $|1.32^{-4} - 1| = 67\%$.

Figure 2: Minimum Quality Standard and Welfare

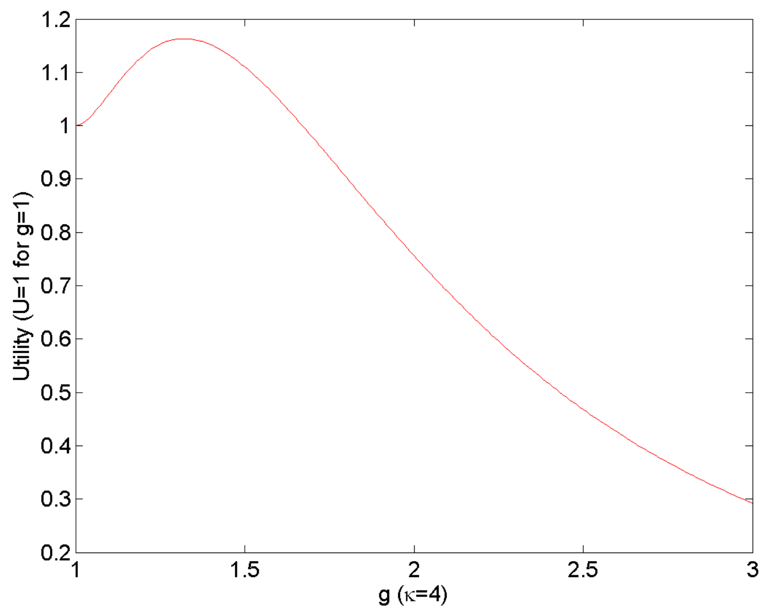

3.1.1 Discussion

Why does the minimum quality standard raise welfare? The standard affects utility through three channels: the number of varieties available for consumption, the average quality across products available, and the business stealing bias. Let us consider the following decomposition of the utility function (2):

$$U = \frac{a^2}{2} \underbrace{N}_{\text{Love for Variety}} \underbrace{\left[1 - \frac{\bar{p}_z}{a}\right]}_{\text{Average Quality}} \underbrace{\frac{\bar{\mu}}{1+\bar{\mu}}}_{\text{Business Stealing Bias}} \tag{16}$$

First, the utility is increasing in the number of varieties available for consumption N. As it is common in models that feature love for variety, consumers' welfare increases with the set of varieties available for consumption. Second, the utility declines with the quality-weighted average price $\bar{p}_z$. Finally, the utility increases with the average markup $\bar{\mu}$ (13) — which represents the business stealing bias. A larger business stealing bias lowers the average markup in the economy, reducing utility.

The number of varieties available for consumption N declines with the quality standard:

$$N = J\left(\frac{ab}{c}\right)^{\kappa} g^{-\kappa} = \frac{L}{f_E}\frac{g^2 + \frac{\kappa}{\kappa+2} - \frac{2\kappa}{\kappa+1}g}{g^2 - \frac{\kappa}{\kappa+2}}\left(\frac{ab}{c}\right)^{\kappa} g^{-\kappa} \tag{17}$$

The mass of entrants J increases with the restriction: by forcing out low-quality firms, the average profits of surviving firms increase, thus, improving the expected profits of potential entrants[16]. However, by raising the quality cutoff, the standard reduces the proportion of entrants which survive to produce. This latter selection effect dominates and N declines with g, therefore reducing welfare all else equal.

Let us consider the average quality-weighted price $\bar{p}_z$:

$$\bar{p}_z = \int_{\bar{z}}^{\infty}\left(\frac{p(z)}{z}\right)\frac{\kappa\bar{z}^{\kappa}}{z^{\kappa+1}}dz = \frac{a}{2g}\left[g + \frac{\kappa}{\kappa+1}\right]$$

The average price is declining in g: raising a country's standard increases the average quality in the economy. Although high-quality firms charge higher markups, their quality-weighted price is lower than low-quality firms. As a result, the quality-weighted price in the economy declines, improving utility.

The reduction in the number of varieties and the increase in the average quality are a standard result of quality standards, which would arise in any model with heterogeneous firms. If preferences were CES, we know that a quality standard must reduce welfare: starting from the efficient allocation, lowering variety dominates the increase in average quality. What drives the welfare improving property of the quality standard is the presence of the business stealing bias, exemplified by $\bar{\mu}$. The sales-weighted average markup $\bar{\mu}$ equals:

$$\bar{\mu} = \frac{\int_{z^*}^{\infty}\mu_{ij}r_{ij}(z)\kappa\frac{b^{\kappa}}{z^{\kappa+1}}dz}{\int_{z^*}^{\infty}r_{ij}(z)\kappa\frac{b^{\kappa}}{z^{\kappa+1}}dz} = \frac{g^2 + \frac{\kappa}{\kappa+2} - \frac{2\kappa}{\kappa+1}g}{g^2 - \frac{\kappa}{\kappa+2}} \tag{18}$$

which is proportional to J, and increases with g. A larger restriction, by forcing out of the market low-quality, low-markup firms, increases the average markup in the economy. Since the average private markup is below the optimal one, the restriction, by increasing such average markup, has a positive effect on welfare. The business stealing bias predicted by the model is negatively related to $\bar{\mu}$: the higher the $\bar{\mu}$ is, the lower the business stealing bias.

A corollary to the fact that the distortion is driven by the business stealing bias is that it depends on the parameter, κ. The "distance" between the market-determined and the

[16]Since the optimal allocation chosen by a planner requires J to be identical to the market allocation, a quality standard, alone, cannot bring the economy to the efficient allocation.

planner's allocation is exemplified by ϵ (12) and $\bar{\mu}$ (13). For large values of κ, the private average markup tends to the socially optimal markup, and the two variables are equal for $\kappa \to \infty$. With symmetric firms ($\kappa \to \infty$), there is no over-/under-production. As we introduce heterogeneity (finite κ), the markup dispersion increases, as does the business stealing bias, which reduces utility as seen in the decomposition (16)[17]. Therefore, the optimal level of the quality standard declines with κ.

3.2 Alternative Preferences

We now extend our framework to several examples with variable markups to provide a better understanding of when a quality standard raises welfare. Under Indirectly Additive (IA) preferences, the relative demand of two goods does not depend on the prices of other goods (Bertoletti and Etro, 2017), but with a non-homothetic functional form, the demand elasticity of the firm is a function of its price relative to total expenditure. We examine two cases in detail, where heterogeneous firms choose non-constant markups due to specific properties of demand: first, non-homothetic separable preferences where firm demand is governed by a general equilibrium object that is a function of both income and the number of competitors, and, second, homothetic and non-separable preferences where the general equilibrium object is governed only by the number of competitors[18]. In summary, we find that the non-homothetic case allows for a welfare improving quality standard although the optimal standard is smaller than our baseline IA case, while the homothetic case allows no room for a quality standard.

3.2.1 Stone-Geary

We start with an example of the additively separable and non-homothetic preferences. These fit into the work of Krugman (1979), with recent variations explored by Behrens et al. (2014), Sauré (2012), Jung et al. (2015), among others. The Stone-Geary case of Simonovska (2015) is especially tractable and we adopt it here. L consumers have the following preferences:

$$U = \int_{\Omega} u(\omega) d\omega = \int_{\Omega} \ln \left[\frac{z(\omega) x^c(\omega)}{\bar{q}} + 1 \right] d\omega \tag{19}$$

[17] In a setting where entry is not exogenous, even with symmetric firms there exists an inefficiency between the level of production and entry (Dixit and Stiglitz, 1977). Since our model produces efficient entry, there is no bias in the symmetric firm case.

[18] Both of these cases fit within the Arkolakis et al. (2017) theoretical framework. We have also examined the Benassy (1996) CES preferences that feature a "variety externality." Depending on the sign of the externality, the market allocates too many or too few varieties, so that in the case of a negative externality it is obvious that a quality standard raises welfare.

where all variables have the same interpretation of the previous case. Moreover, let us spare some notation by setting wages and per capita income equal to one. The model details can be found in appendix 8.2, here we jump straight to a decomposition of utility[19]:

$$U = N\bar{u} = \frac{N}{2}\left[\frac{1}{\kappa} + \ln g\right] \tag{20}$$

$$= \frac{(L\bar{q}cb^{\kappa})^{\frac{1}{\kappa+1}}}{2\bar{q}cf_E^{\frac{1}{\kappa+1}}} \frac{\left[g - \frac{\kappa}{\kappa+0.5}g^{0.5} + \frac{\kappa}{\kappa+1}\right]^{\frac{1}{\kappa+1}}}{g - \frac{\kappa}{\kappa+0.5}g^{0.5}} \left[\frac{1}{\kappa} + \ln g\right], \tag{21}$$

where $\bar{u}$ is the average quality. In this case, love for variety and business stealing bias are both represented within N. As in the Addilog case, there is a level of $g > 1$ that maximizes welfare. In the appendix, we report utility for $g \geq 1$ and Figure 7 shows that a small quality requirement improves the utility of consumers.

The intuition is similar to our baseline case. The market allocation allows for *too many* varieties: the business stealing bias that was present in the baseline model is a feature of the Stone-Geary case as well. Figure 3 compares the optimal quality standard for the Stone-Geary case to the IA case. The optimal g is greater than one for both models, although it is larger in the IA case. This is due to the fact that with separable preferences, a reduction in N simultaneously reduces the"toughness" of competition. Less competition allows high-quality firms to raise their markup relatively more than low-quality firms, thus reallocating production to low-quality firms.

[19]We do note the all-important cutoff: $z^{*(SG)} = \bar{q}c\lambda^{SG}$, which depends on a general equilibrium variable — the marginal utility of income (λ^{SG}). As a result, markups do not just depend on real income, but also on the number of varieties available for consumption.

Figure 3: Optimal Minimum Requirement: Stone-Geary VS Indirectly Additive

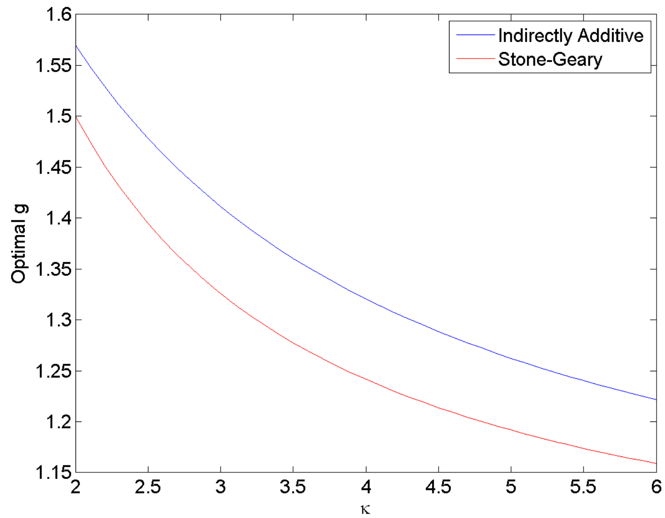

3.2.2 QMOR

Next, we examine the "Quadratic Mean of Order R," a symmetric expenditure function that generates homothetic demand and variable markups. The general expenditure function is detailed in Feenstra (2014), and here we take a special case with linear demand (assuming $r = 2$ in that paper). There is no closed form for utility, so we follow Arkolakis et al. (2017) to write the demand as a function of general equilibrium objects, adding a measure of quality as in the previous cases:

$$x(z) = \gamma Q \left(\frac{p(z)}{zP} - 1 \right), \tag{22}$$

where P is the choke price as defined in Feenstra and Weinstein (2016), γ is a parameter, and Q is a demand shifter which for now we treat as a constant because firms take it as given[20]. We can solve for the firm price: $p(z) = \frac{1}{2}c\left(\frac{z}{z^{*(Q)}} + 1\right)$, with $z^{*(Q)} = \frac{c}{P}$. As shown in Feenstra and Weinstein (2016), the general equilibrium object in this case, P, depends only on the number of competitors[21].

The details of this model are in appendix 8.3. Although we do not have a social planner allocation in this case, we can derive welfare for a given g (defined as before). Notice that

[20] Q is defined such that the budget constraint holds: $\int_\Omega \gamma p(z) Q \left(\frac{p(z)}{P} - 1\right) dz = y$.

[21] $P = \left[\frac{N}{N - (\tilde{N} + \gamma/\beta)}\right] \int_{z^*}^{\infty} \frac{1}{N} p(z) dz$. β is another parameter which will not play a role, and $\tilde{N}$ is the set of all *possible* varieties.

welfare is the inverse of expenditure function, which turns out to be:

$$e_2(\mathbf{p}) = \left(\frac{1}{f_e}\right)^{1/2} \left[\frac{\left(-\frac{1}{\kappa-1} g^{\kappa-1} - \frac{2}{\kappa} g^{\kappa} - \frac{1}{\kappa+1} g^{\kappa+1}\right)}{\left(\frac{1}{-\kappa+1} g^{\kappa-1} + \frac{1}{\kappa+1} g^{\kappa+1}\right)} \right], \tag{23}$$

This expression can be analyzed numerically, and it increases with g. In the QMOR, the selection mechanism is *too strong* as it drives out too many firms. Any binding minimum quality standard further *reduces* welfare.

4 Quality Standards and Trade Openness

A quality standard improves welfare in a closed economy when the market allocation generates business stealing bias, and too many low-quality firms are active in the market. Such an outcome occurs when firms are heterogeneous and monopolistically competitive, and when consumers exhibit non-homothetic preferences. The purpose of this section is to study whether international trade modifies the optimal level of a quality requirement. Given the secular decline in tariffs and traditional non-tariff measures, technical measures and minimum quality requirement, are subject of debates over trade agreements (Baldwin et al., 2000), and of several concerns raised at the WTO (Fontagné et al., 2015).

To address the policy debates over minimum quality standards, we consider a world made of two identical economies, in which exports require an iceberg trade cost τ. As the two countries are identical, we can normalize per capita income and wages to one in both countries. We assume that the two countries decide a common minimum quality requirement $\bar{g}$, which holds in each nation, and we study how different levels of the iceberg trade cost τ affect the optimal level of $\bar{g}$. Thus, we abstract from political economy issues — as well as a complicated non-cooperative game — that would arise if each country were to choose its own optimal level of $\bar{g}$.

We can interpret our thought experiment as the problem of an international authority (such as the European Union) or of a negotiation process between two countries deciding a common minimum quality requirement for a differentiated good. The optimal requirement depends on the level of openness of the two countries, which is controlled by τ. Since the two countries are identical, the optimal requirement does not depend on different bargaining powers between two countries with different incentives.

Similarly to the closed economy case, let $\bar{z} = \bar{g}$ denote the firms with quality z equal to the quality standard. Moroever, let $g = \frac{\bar{z}}{z_D^*}$ be a measure of the restrictiveness of the requirement, where z_D^* is the domestic quality cutoff (8). As in a standard Melitz (2003) model, the presence of iceberg trade costs divides firms in exporters and non-exporters. Only high-quality firms are exporters, since their quality level is large enough to sell their goods abroad. As a result, the effect of g on the selection of firms exhibit a discontinuity. In particular, for $g \in [1, \tau]$, the quality standard only affects low-quality non-exporters. On the other hand, for $g \in [\tau, \infty)$, the standard is so restrictive that the only firms able to remain active are high-quality exporters. We leave the derivations to the appendix, and report here the utility function of the representative consumers of both countries:

$$U^T(g) = \left(\frac{ab}{c}\right)^{\kappa} \frac{a^2 L}{8 f_E} \begin{cases} \dfrac{\left[g^{-\kappa} + \frac{\kappa}{\kappa+2} g^{-2-\kappa} - \frac{2\kappa}{\kappa+1} g^{-1-\kappa} + \frac{2\tau^{-\kappa}}{(\kappa+1)(\kappa+2)}\right]^2}{g^{-\kappa} - \frac{\kappa}{\kappa+2} g^{-2-\kappa} + \frac{2\tau^{-\kappa}}{\kappa+2}} & \text{if } g \in [1, \tau] \\ g^{-\kappa} \dfrac{\left[2 + \frac{\kappa(1+\tau^2)}{\kappa+2} g^{-2} - \frac{2\kappa(1+\tau)}{\kappa+1} g^{-1}\right]^2}{2 - \frac{\kappa(1+\tau^2)}{\kappa+2} g^{-2}} & \text{if } g \in [\tau, \infty) \end{cases} \tag{24}$$

$U^T(g)$ is continuous and, as shown in figure 4, is non-monotone and hump shaped. Hence, even in the presence of trade, a minimum quality requirement is welfare improving.

Figure 4: Minimum Requirement and Welfare

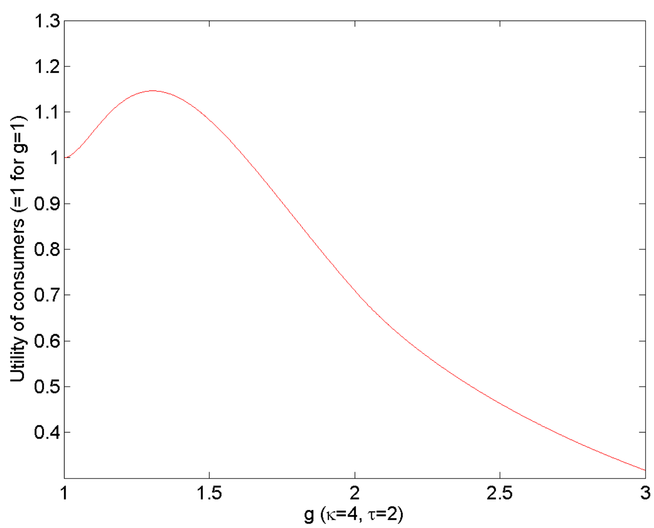

How does the level of international openness affect the optimal minimum quality requirement? To answer to this question, we compute the optimal $\bar{g}^{opt}$ for different levels of τ. As shown in figure 5, lower trade costs reduce the optimal quality requirement[22].

[22] We restrict τ so that $g^{opt} \in [1, \tau]$, and the restriction only affects non-exporters. Although it is theoretically possible for the standard to force out of the market all non-exporters, it is not a realistic outcome and, thus, we ignore it.

Figure 5: Optimal Requirement and Trade Costs

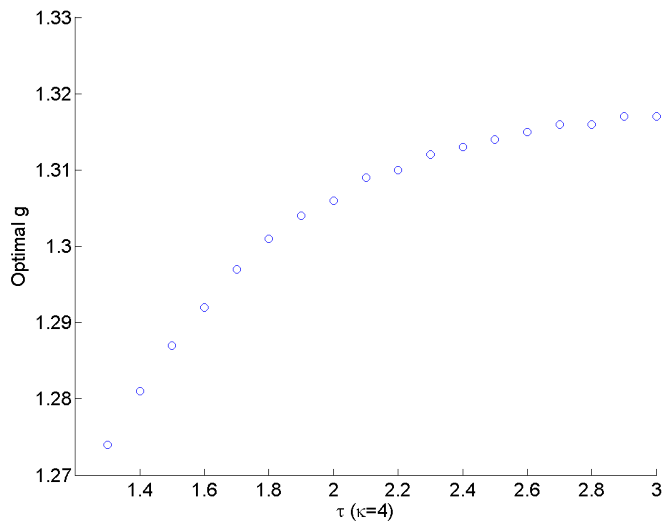

Under IA preferences, the domestic cutoff $z_D^* = \frac{c}{a}$ is not affected by trade costs, nor is the mass of entrants J, which, under no quality restriction, is identical to the baseline model. Hence, without the quality restriction, international trade does not bring the economy closer to the outcome chosen by a planner along those two dimensions. However, it is reassuring that even in the extreme case where there is no selection effect that drives domestic firms out of the market in response to a trade liberalization (due to the assumption of IA preferences), lower iceberg trade costs still improve allocative efficiency.

The rationalization of the above finding comes through market share reallocation. With lower iceberg trade, the ratio of quantities chosen by the planner relative to that determined by the market, is closer to one. Since the market allocation causes an overproduction of low-quality goods and an underproduction of high-quality goods, it is the exporters' varieties — which are of high quality — that are under-produced relatively to domestic low-quality varieties. One way to reduce such a distortion is for the authority to force the low-quality varieties out of the market. In fact, that is exactly the mechanism that we studied in the previous sections. Lower variable trade costs improve allocative efficiency through a different mechanism: as trade costs decline, export production increases. Although there are the same number of domestic varieties available, these now compete with foreign varieties, which reallocates market share towards high-quality goods. Therefore, the level of distortions is *lower* in a more open economy, alleviating the need for a minimum quality requirement. We believe this rationalizes a dual approach for policymakers: pushing towards lower trade costs while lowering the restrictiveness of quality standards.

5 Data

A. Detailed Database of Non-Tariff Measures

In order to take the model predictions on quality standards to the data, we make use of the prevalence of non-tariff measures (NTMs). NTMs are an important way through which local policy makers can limit market access through the imposition of regulatory measures. With the significant decline in import tariffs, trade economists have pointed towards non-tariff technical measures as an increasingly relevant subject in trade agreements (Maskus et al., 2000; Baldwin et al., 2000). Specifically, sanitary and phytosanitary (SPS) and technical barriers to trade (TBT) are now the most crucial non-tariff policies (UNCTAD, 2017). NTMs have been considered as a form of trade costs that primarily impact the extensive margin (Ferro et al., 2015; Fontagné et al., 2015), and therefore provide a consistent mapping to the quality standard explored in section 3.

A new database of NTMs has recently been made available through TRAINS which provides comprehensive information of measures imposed by WTO members.[23] It builds on the Specific Trade Concerns (STC) data utilized in Fontagné et al. (2015) — which lists concerns raised by trading partners — to include *all* domestic regulations found in official texts that can be classified as a NTMs. For our purposes, it is crucial that these regulatory measures affect importers and domestic firms alike, since our model makes predictions about domestic production.

As described in UNCTAD (2017): " transparency and regulatory convergence can substantially reduce trade costs while allowing countries to pursue policy objectives...the objective of the TRAINS database is to increase the transparency and understanding about regulations and trade control measures." The 2012 NTM classification separates measures into 16 chapters (labeled A-P), and we make use of the first three chapters: SPS TBT, and preshipment inspections, to construct our measure of quality regulation. These three chapters are "technical measures," and therefore fit most closely with our quality standard in the theory. We should note however that the vast majority of listed NTMs are either SPS or TBT. An important advantage of this data for our purposes is that the technical regulations apply to both imported products and locally-produced goods.

TRAINS collects official measures imposed by countries that might affect international trade, that are mandatory, and that are currently applied. National governments or local consultants hired by the World Bank collect regulations from official government sources, such as Customs Agencies or Government Ministries. Within a regulation, each imposed measure is taken into account and all the affected products (6 digit HS classification) and

[23]The web application to retrieve the data is available at http://i-tip.unctad.org/.

countries are listed.[24] Although each measure lists the year the measure was first introduced, we do not believe it makes sense to analyze variation of regulation with a time dimension, so we consider the total number of measures in force applied to each product at the time the data was constructed (2015).

From the reporter (imposing country)-product level NTM data, we must aggregate the data in order to merge it to our domestic production data. For the most aggregated analysis we use the NTM-MAP, a CEPII database described in Gourdon (2014). This database measures the incidence of NTM measures in each reporting country at the HS 2-digit level of aggregation. We also construct our own incidence measure at a more disaggregate level in order to merge the data to firm production data. Firms are classified by 4-digit ISIC (revision 3) codes, so we construct the incidence of NTM measures for each ISIC code by aggregating the HS 6-digit product information.

B. Eurostat Business Demography

In order to test the model predictions on entry and exit we use Eurostat's database of business demography. For a subset of EU countries that we can merge to the NTM-MAP database, the number of factories, death rate, survival rate, and birth rate are provided for various size classes at the NACE industry level.[25] We distinguish between "small firms," those with fewer than 10 employees, and "large" firms with over 10 employees. We believe that these two size classes can be mapped to the model such that more regulated industries should experience higher entry and exit rates of "low-quality," or small, firms. We manually concord the HS 2-digit classification of NTM measures to the NACE (revision 2) industry data (of which we use only manufacturing industries). Since we will not be examining the time dimension, we can pool the annual data by either choosing only one year of data or taking the average of industry-country observations across all observed years. The fact that this industry-level data is provided for multiple countries will allow us to control for industry and country fixed effects. [SUMMARY STATS: NACE INDUSTRY ROWS WITH AVERAGE ENTRY AND EXIT RATES PLUS NTM MEASURES AS COLUMNS?]

C. Chilean Firm Data

Finally, we would also like to test the distributional predictions made by the model that are the consequences of quality standard measures. To do so, we construct distributional predictions, comparing high and low quality firms across industries differentiated by their

[24]We only consider measures imposed on all countries, and not on individual trade partners.

[25]More information on this database available at: http://ec.europa.eu/eurostat/web/structural-business-statistics/entrepreneurship/business-demography.

level of regulation. [WILL PUT MORE INFO ON THE FIRM DATA, CAN JUST TAKE FROM MY JMP.]

6 Empirical Analysis

In section 3 we show that a quality standard improves allocative efficiency in a model with monopolistically competitive firms and non-homothetic demand. In this section we test whether the results in that model are indeed supported by the data. The first step is to establish that regulations can in fact raise the cutoff to produce in the domestic economy, the empirical counterpart to (15). [HERE IT IS UNSURE WHETHER WE USE "DOMESTIC" (HOPEFULLY) OR "EXPORTERS", OR BOTH]. To this end, we test whether regulations in the form of sanitary & phytosanitary (SPS) and technical barriers to trade (TBT) imposed by the host country makes survival less likely and exit more likely. Second, to test the validity of the welfare results summarized by (16), we construct measures that can proxy allocative efficiency. We then investigate whether regulations that raise quality standards can indeed reduce the misallocation bias highlighted in the theory.

6.1 Entry and Exit by Size Class

There are various entry and exit predictions salient to our welfare implications of the quality standard that we take to the data. Recall that the total number of varieties produced in i with a quality standard g, is:

$$N = J(g)\left(\frac{b}{g}\right)^{\kappa}.$$

We consider several comparative statics that are derived from section 3. We have shown that $\frac{dN}{dg} < 0$, since the number of varieties is smaller with a higher standard. Furthermore, it can be shown that $\frac{dJ}{dg} > 0$, so that the number of potential entrants *increases* with the quality standard. The intuition is straightforward: there are fewer active firms in equilibrium, raising average profits and therefore incentivizing more firms to try to enter. A higher number of potential entrants and a lower number of available varieties then implies that a quality standard *lowers* the survival rate and *raises* the death rate, two measures that are available from the Eurostat demographics database.

The comparative statics above can also be applied differentially across firms. From Figure 1 it is clear that the quality standard eliminates the lowest quality firms. Therefore, the survival and death rate implications above should *be stronger* for low quality firms, which in our model is synonymous with size. We therefore run the following specification using the Eurostat data:

$$y_{sci} = \alpha + \alpha_c + \alpha_i + \gamma_E NTM_{ci} + \delta_E Size_s + \beta_E NTM_{ci} * Size_s + \epsilon_{sci}, \tag{25}$$

where s is the size class – "small" or "large" – c refers to a country, and i the sector. Our outcome variables will be the number of active firms, the birth rate, and the death rate. All the outcome measures are available by size class.[26] In order to control for country-industry characteristics, we normalize the outcome of each size class using the outcome variable for the country-industry as a whole.[27] Since we can only normalize the outcome variables by a country-industry measure in the case with size classes, we believe this is the best specification. Standards, or our measure of NTMs in the data, vary at the country-sector level. Therefore, we control for country and sector fixed effects.[28] 'NTM" refers to non-tariff measures, which we proxy with a "prevalence" index of SPS and TBT measures imposed. This variable maps to $\bar{g}$ in the theoretical model.

We run each specification in two ways: pooliing the size classes together (in this case there is no size dummy), and also by size class (the main specification). In the latter case we are interested in the coefficient β_E, which is the differential effect of the outcome variable between small and large firms.

As stated above, we concord the Eurostat data with NTM measures at the 2-digit HS classification available in CEPII. The NTM measures are not available for al years, so we create one regulation measure with no time variation. The Eurostat outcome measures are mostly available annually going back to 2006, so we can run specification 25 in two ways. First, we average each outcome variable across all years to eliminate any time variation. This way we are assured that the results are not driven by macroeconomic shocks that are concurrent to the years of data that we use. Second, to make use of a larger sample size, we use all the size-country-sector-time data, but control for country-year and sector-year fixed effects. In this way we are once again only using within sector and within country variation, and not capturing any time variation in the outcome measures.

Table 1 displays the results. In Panel A we average the business demography measures over all years. The population of active firms do not change significantly with the NTM

[26] The Eurostat data classifies firms by "1-4, 5-9, and 10+ employees." We call the first two classes "small" and the latter class "large."

[27] For example, we divide the birth rate of small firms in the Textile sector in France by the total birth rate in the Textile sector in France.

[28] We do not have country-sector fixed effects in the regression with size classes because we are already normalizing each size class outcome with the total outcome at the country-industry level.

measures. However, both entry rates and death rates behave as we would expect. The total entry rate is higher for a higher prevalence of NTMs (column (3)), and although this is not significant, column (4) does show that the entry rate is significantly larger for small firms. This is consistent with the model, where a quality standard raises expected profits and increases the number of potential entrants. The rise in potential entrants should come from smaller, or low-quality firms. Columns (5) and (6) show similar results for the death rate. Once again, a higher prevalence measure increases the death rate, and it is significantly larger for smaller firms. As expected, the quality standard makes it more difficult for small, low-quality firms to survive. Panel B reproduces the same specification but using all years of data with country-year and sector-year fixed effects.

Table 1: Entry and Exit by size class in response to NTMs

Panel A: Pooled across Years

	Population of Active Firms		Entry Rate		Death Rate	
	(1)	(2)	(3)	(4)	(5)	(6)
Prevalence Score	-0.119	-0.026	0.345	-0.033	0.252	0.247
	(0.230)	(0.062)	(0.341)	(0.180)	(0.157)	(0.151)
small		0.194***		0.360***		0.230***
		(0.054)		(0.054)		(0.037)
Score*Small		-0.044		0.054*		0.075***
		(0.036)		(0.028)		(0.022)
R^2	0.703	0.542	0.376	0.665	0.511	0.592
# Observations	339	678	323	645	320	642

Panel B: All Years

	Population of Active Firms		Entry Rate		Death Rate	
	(1)	(2)	(3)	(4)	(5)	(6)
Prevalence Score	-0.027	0.021	0.155	0.102	-0.062	-0.073
	(0.045)	(0.024)	(0.157)	(0.095)	(0.246)	(0.123)
small		0.183***		0.391***		0.283***
		(0.053)		(0.042)		(0.022)
Score*Small		-0.036		0.042*		0.053***
		(0.032)		(0.023)		(0.002)
Fixed Effects	C-T, I-T	C-T, I-T	C-T, I-T	C-T, I-T	C-T, I-T	C-T, I-T
R^2	1	1	0	1	1	1
# Observations	2637	5274	2490	4983	2416	4835

.

6.2 Firm Distributions

Aside from entry and exit dynamics, the welfare results are reliant on the fact that quality standards alter the sales distribution across quality. A quality standard below the optimal standard raises welfare by taking market share from a low-quality firm and reallocating it to a higher-quality firm. As long as the reduction in the business-stealing bias and increase in average quality are larger than the welfare loss due to less variety, then welfare increases.

In the next set of results, we use firm-level data to test whether non-tariff measures indeed reallocate market share from low-quality to high-quality firms. To do so, we use a census data of Chilean firms that reports revenues and a variety of input measures. The input measures are used to construct a proxy for quality: we take the firms' capital stock, labor costs, and intermediate input costs, and divide each by the number of employees. Higher capital intensity, wages per worker, and input costs all arguably correlate with quality and have been used in previous studies (Kugler and Verhoogen, 2012; Hallak and Sivadasan, 2013). In order to create a quality indicator, we label a firm as "high-quality" if it is above the median in the quality proxy within its industry. In addition, we will use an indicator for whether the firm is an exporter *or* importer as a measure of quality.

Firms report their industry at the 4-digit ISIC level of detail. Therefore, we can use the product level measure of NTMs provided by TRAINS and aggregate them to get the same level of aggregation. We construct a prevalence measure of NTMs in each industry by counting the total number of measures imposed in an industry relative to the number of HS 6 products in a 4-digit ISIC industry. The TRAINS data provides the year that each NTM is imposed. For that reason, we can track the number of measures imposed by year to leverage time variation in regulations. However, it is not clear that firms will react to the imposition of the measures during the year they are imposed. For that reason, in our main analysis we create a total prevalence measure for NTMs that aggregates all measures in force in 2014 (the final year of data), as in the CEPII data we use above. The regression specification will be the following:

$$sales_{fi} = \alpha + \alpha_i + \gamma_M NTM_i + \delta_M Quality_f + \beta_M NTM_i * Quality_f + \epsilon_{fi}, \qquad (26)$$

where not f is the firm and i continues to indicate sectors. Notice that the "Quality" indicator is a dummy so that this indicator is not correlated with the level of sales. We cannot control for firm fixed effects in this specification, but we do have sector fixed effects. The coefficient of interest is β_M, which can be interpreted as the difference in the average sales of high-quality firms relative to the average sales of low-quality firms at different levels of NTMs. The model predicts that a higher standard results in relatively more sales of high-quality firms, so that $beta_M$ should be positive.

Table 2 displays the results for different measures of the quality proxy (Capital, Wages, Input costs, and exporter or importer). The firm data is a panel from 1995-2007, but the variation in the NTMs is only at the industry level. Instead of averaging firm sales over all years, we use all firm-year observations and include sector-year fixed effects so that all of the variation is within sectors. Across our proxies for quality, the results are robust with $beta_M$¿0. The ratio of the average sales of high-quality firms relative to the average sales of low-quality firms is larger in industries that have more NTMs. This is consistent with NTMs raising allocative efficiency by reallocating market share away from low-quality firms, at least if this effect is not dominated by the reduction in variety (if the standard were above the optimal).

Table 2: Firm Sales Heterogeneity - by NTMs in Industry

	Domestic Sales			
	(1)	(2)	(3)	(4)
Quality Ind (Capital/L)	0.913***			
	(0.144)			
NTM*Quality Ind (Capital/L)	0.110**			
	(0.046)			
Quality Ind. (Wage/L)		1.439***		
		(0.158)		
NTM*Quality Ind. (Wage/L)		0.052		
		(0.046)		
Quality Ind. (InputValue/L)			1.390***	
			(0.171)	
NTM*Quality Ind. (InputValue/L)			0.123**	
			(0.052)	
Exp/Imp				1.747***
				(0.107)
NTM*Exp/Imp				0.084*
				(0.045)
Fixed Effects	Industry-Year	Industry-Year	Industry-Year	Industry-Year
R^2	0.357	0.442	0.463	0.444
# Observations	63840	63840	63840	63840

The measure of NTMs is at the 4 digit ISIC industry level. The total number of measures in each industry are summed over all years to create one regulation index for each industry. We include industry fixed effects and cluster the standard errors at the industry level.

The results above are problematic if they are driven by firm characteristics that are correlated with the regulation restrictions in the sector. For example, the regulations are aggregated from the HS 6 product level, so that firms within the same 4-digit ISIC might actually be exposed to different levels of regulation. Although firms only report their 4-digit industry, firms within these industries could be producing different products at a more disaggregate level. The only way to convincingly control for firm-specific characteristics is to have time variation in the imposition of standards. Therefore, as a robustness check we extend the specification above to also allow for time variation in NTMs.

The TRAINS data does report the year that each measure was imposed. We use this information to count the number of measures imposed in each industry during the firm panel range, 1995-2007. In this case run the following specification:

$$sales_{fit} = \alpha + \alpha_{it} + \alpha_f + \gamma_M NTM_{it} + \delta_M Quality_f + \beta_M NTM_{it} * Quality_f + \epsilon_{fit}. \quad (27)$$

The results are reported in Table 3. It is reassuring that the β_M is still mostly positive. Although the coefficient drops in magnitude, using capital intensity and wages per worker as quality indicators still yields differential effects across industries that are significant at the 5% level. Although not significant, the results are fairly strong for the input quality indicator as well. Only when using importer/exporter as a quality indicator does there not seem to be any distributional effects from the NTMs. Overall, these results indicate that the specification without firm fixed effects is robust and not contaminated by unobserved firm characteristics.

Table 3: Firm Sales Heterogeneity - by NTMs in Industry-Year

	Domestic Sales			
	(1)	(2)	(3)	(4)
NTM*Quality Ind (Capital/L)	0.039**			
	(0.016)			
NTM*Quality Ind. (Wage/L)		0.038**		
		(0.019)		
NTM*Quality Ind. (InputValue/L)			0.023	
			(0.015)	
NTM*Exp/Imp				0.001
				(0.019)
Fixed Effects	Firm, I-Y	Firm, I-Y	Firm, I-Y	Firm, I-Y
R^2	0.957	0.957	0.957	0.957
# Observations	41236	41236	40787	41236

The measure of NTMs is at the 4 digit ISIC industry level. The NTM measure is computed by year. We include industry and firm fixed effects and cluster the standard errors at the industry level.

6.3 Quantitative Analysis

7 Conclusion

References

C. Arkolakis, A. Costinot, and A. Rodriguez-Clare. New trade models, same old gains? *American Economic Review*, 102(1):94–130, 2012.

C. Arkolakis, A. Costinot, D. Donaldson, and A. Rodrguez-Clare. The Elusive Pro-Competitive Effects of Trade. *NBER Working Paper*, (21370), 2015.

C. Arkolakis, A. Costinot, D. Donaldson, and A. Rodrigues-Clare. The elusive pro-competitive effects of trade. Working paper, 2017.

R. E. Baldwin, J. McLaren, and A. Panagariya. Regulatory protectionism, developing nations, and a two-tier world trade system [with comments and discussion]. *Brookings Trade Forum*, pages 237–293, 2000.

P. Bastos and J. Silva. The Quality of a Firm's Exports: Where you Export Matters. *Journal of International Economics*, 82(2):99–111, 2010.

S. Basu and J. Fernald. Aggregate productivity and aggregate technology. *European Economic Review*, 46:963–991, 2002.

K. Behrens, G. Mion, J. Suedekum, and Y. Murata. Trade, Wages, and Productivity. *International Economic Review*, 55(4):13051348, 2014.

J.-P. Benassy. Taste for variety and optimum production patterns in monopolistic competition. *Economics Letters*, 52(1):41–47, 1996.

P. Bertoletti and F. Etro. Monopolistic competition when income matters. *Economic Journal*, In Press, 2017.

P. Bertoletti, F. Etro, and I. Simonovska. International trade with indirect additivity. *American Economic Journal: Micro*, In Press, 2017.

T. Chaney. Distorted Gravity: The Intensive and Extensive Margins of International Trade. *American Economic Review*, 98(4):1707–21, 2008.

N. Chen and D. Novy. Gravity, trade integration, and heterogeneity across industries. *Journal of International Economics*, 85(2):206 – 221, 2011.

S. Demidova and A. Rodriguez-Clare. Trade policy under firm-heterogeneity in a small economy. *Journal of International Economics*, 78:100 112, 2009.

S. Dhingra and J. Morrow. Monopolistic Competition and Optimum Product Diversity Under Firm Heterogeneity. 2016.

J. Dingel. Determinants of Quality Specialization. *Mimeo*, 2015.

A. Dixit and J. Stiglitz. Monopolistic competition and optimum product diversity. *American Economic Review*, 67(3):297–308, 1977.

C. Edmond, V. Midrigan, and D. Xu. Competition, markups and the gains from international trade. *American Economic Review*, 105(10):3183–3221, 2015.

R. C. Feenstra. Restoring the product variety and pro-competitive gains from trade with heterogeneous firms and bounded productivity. *NBER Working Paper*, (19833), 2014.

R. C. Feenstra and H. L. Kee. Export variety and country productivity: Estimating the monopolistic competition model with endogenous productivity. *Journal of International Economics*, 74(2):500–518, 2008.

R. C. Feenstra and J. Romalis. International prices and endogenous quality. *The Quarterly Journal of Economics*, 129(2):477–527, 2014.

R. C. Feenstra and D. Weinstein. Globalization, competition, and u.s. welfare. *Forthcoming in Journal of Political Economy*, 2016.

G. Felbermayr, B. Jung, and M. Larch. Optimal tariffs, retaliation, and the welfare loss from tariff wars in the melitz model. *Journal of International Economics*, 89(1):13–25, 2013.

A. M. Fernandes, E. Ferro, and J. M. Wilson. Product standards and firms? export decisions. *World Bank Working Paper Series*, (7315), June 2015.

E. Ferro, T. Otsuki, and J. S. Wilson. The effect of product standards on agricultural exports. *Food Policy*, 50(C):68–79, 2015.

L. Fontagné, G. Orefice, R. Piermartini, and N. Rocha. Product standards and margins of trade: Firm-level evidence. *Journal of International Economics*, 97(1):29–44, 2015.

J. Gourdon. CEPII NTM-MAP: A Tool for Assessing the Economic Impact of Non-Tariff Measures. *CEPII Working Papers*, (2014-24), Dec. 2014.

J. C. Hallak and J. Sivadasan. Product and process productivity: Implications for quality choice and conditional exporter premia. *Journal of International Economics*, 91(1):53–67, 2013.

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

J.-W. Jung, I. Simonovska, and A. Weinberger. Exporter heterogeneity and price discrimination: A quantitative view. *NBER Working Paper*, (21408), 2015.

P. Krugman. Scale economies, product differentiation, and the pattern of trade. *The American Economic Review*, 1980.

P. R. Krugman. Increasing returns, monopolistic competition, and international trade. *Journal of international Economics*, 9(4):469–479, 1979.

M. Kugler and E. Verhoogen. Price, plant size, and product quality. *The Review of Economic Studies*, 79:307–339, 2012.

G. Mankiw and M. Whinston. Free Entry and Social Inefficiency. *The RAND Journal of Economics*, 17(1):48–58, 1986.

K. Manova and Z. Zhang. Export Prices Across Firms and Destinations. *Quarterly Journal of Economics*, 127(1):379–436, 2012.

J. Martin. Markups, Quality, and Transport Costs. *European Economic Review*, 56(4): 777–791, 2012.

K. E. Maskus, J. S. Wilson, and T. Otsuki. Quantifying the impact of technical barriers to trade : a framework for analysis. (2512), 2000.

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

D. Restuccia and R. Rogerson. Policy Distortions and Aggregate Productivity with Heterogeneous Plants. *Review of Economic Dynamics*, 11(4):707–720, 2008.

P. Sauré. Bounded Love of Variety and Patterns of Trade. *Open Economies Review*, 23(4): 645–674, 2012.

I. Simonovska. Income differences and prices of tradables: Insights from an online retailer. *The Review of Economic Studies*, 82(4):1612–1656, 2015.

UNCTAD. International classification of non-tariff measures. 2012.

UNCTAD. UNCTAD Trains: The Global Database on Non-tariff Measures, Version 2.0 2016. Working paper, United Nations, 2017.

X. Vives. *Oligopoly Pricing: Old Ideas and New Tools*, volume 1 of *MIT Press Books*. The MIT Press, 2001.

A. Weinberger. Markups and Misallocation with Evidence from an Exchange Rate Appreciation. Working paper, 2017. URL http://aweinberger.weebly.com/uploads/2/9/2/5/29251877/markups_misallocation_d3.pdf.

8 Appendix

8.1 Addilog Model

Figure 6: Minimum Quality Standard and Quality Dispersion

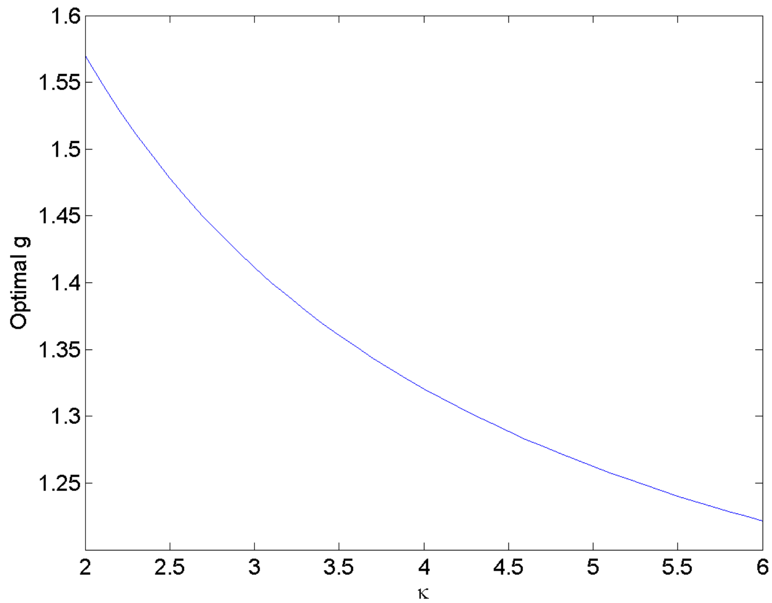

8.2 Stone-Geary Model

Figure 7: Minimum Requirement and Welfare

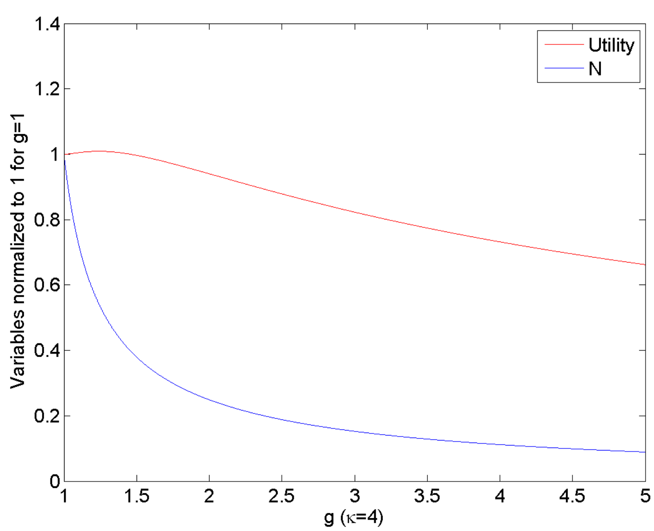

8.3 QMOR Model