Justin Clifford
NetID: jcliffo7
ECO 389W
Spring 2020

The Purchasing Power Parity Puzzle in the Context of Strategic Transport Markets[1]

Submitted to the Department of Economics to Fulfill the Requirements for the Degree of Bachelor of Arts, With Honors

Abstract

I build two formulations of a 7-country algorithmic model of global trade flows to simulate relative price levels between each pair of these seven countries; if this simulated relative price level (SRPL) is equivalent to nominal exchange rate (NER) between those two countries at a given time, Purchasing Power Parity (PPP) conditions are satisfied. First, I consider a more conventional version of the model in which "Iceberg" costs are the only contributor to trade frictions among producers. Then, I introduce a representative profit-maximizing shipping firm to the model. Adding such a firm decreases R^2 between SRPL and NER for 17 out of 21 country pairings studied; and overall, I find that profit-maximizing behavior in shipping markets accounts for approximately 8% of deviation from PPP conditions. As an ancillary finding, the model predicts that every 1% increase in shipping price markups makes the satisfaction of PPP conditions 0.5575 percentage points less likely.

[1] I am grateful to Professor Michael Rizzo and my ECO 389 classmates for the support and guidance they have given me while working on this paper, especially in light of the ongoing COVID-19 crisis. In addition, I would like to thank Professors Dan Lu and Yan Bai, whose courses at UR greatly enhanced my understanding of, and interest in, international economics and trade. I also owe a big thanks to Eytan Buchman and the folks at Freightos, who allowed me to use their Baltic Index data free of charge; further, I am grateful to UR librarians Lydia Auteritano and Kathy Wu for helping me obtain the Freightos Baltic Index dataset. All errors are my own.

1: Introduction

Purchasing Power Parity (PPP) refers to the idea that all currencies should be equally powerful. In other words, when controlling for nominal exchange rates, a diversified basket of goods should have the same price in every country, no matter the currency in which the price of that basket of goods is denominated. If this condition were not satisfied, it would lead to arbitrage opportunities; for if PPP fails between some country *A* and another country *B*, one would be able to buy a basket of goods in *A* using *A*'s currency, and sell it for an immediate profit in *B* using *B*'s currency, or vice versa. Conventional economics would argue that such opportunities should never be available. But PPP often fails in the real world, except in the very long term, and such failures are very well documented in the data (Murray and Papell 2005; many others). And high observed short-term volatility of real exchange rates seems to be inconsistent with slowness of overall adjustments back to PPP conditions (Rogoff 1996); this inconsistency has come to be known as the Purchasing Power Parity Puzzle. So what could be behind this apparent violation of basic economic principles?

Trade costs are thought to be a key perpetrator in failures of PPP, or at least slowness of adjustments back to PPP conditions (Obstfield and Rogoff 2000). This argument makes sense given the example from the previous paragraph; it costs something to transport a basket of goods between *A* and *B*, and this sort of cost could eliminate the arbitrage I otherwise would have gained from such a transaction. However, there are many unique, individual costs, incurred by a variety of actors, wrapped up in the term "trade costs", and conventional PPP models fail to differentiate between the specific difficulties that each kind of trade cost presents.

Which one of these individual costs causes the largest disruption in PPP conditions? Conventional models usually only consider two forms of trade cost: (a) the proportional value that goods lose while enroute to their destination market, often called "Iceberg Costs" or the "Melting Effect", as well as (b) tariffs and quotas imposed by governments. But these models skip over the costs of performing the shipping itself, which are incurred most directly by profit-maximizing shipping firms. These price-making firms choose their rates rationally, as higher rates decrease shipment volumes *ceteris paribus*, and vice versa. Thus, the determination of trade costs is not parametric —as conventional models would have us believe— but rather strategic; for an exporting firm's pricing behavior affects a shipping line's pricing behavior, and vice versa. In addition to the reasons just stated, I suspect that PPP failure has something to do with transport firms and their pricing choices rather than the Melting Effect, given the aforementioned example of countries *A* and *B*: the reason arbitrage opportunities exist in that example is because of the cost of physically transporting the basket of goods from place to place, *not* the value that that basket of goods loses while enroute. The goal of this paper is to endogenize this sort of strategic decision-making that occurs within transport markets. To do so, I build a time-sensitive, algorithmic model of global trade flows. I simulate two cases within this framework: first, I simulate a counterfactual world in which the exogenous Melting Effect accounts for all trade costs, i.e., firms' exporting decisions are simply parametric. I then build a full model with transport markets operating strategically. This model predicts that PPP conditions are more likely to occur in the first case than the second. The conclusion is that profit-maximizing behavior among shipping firms —a usual omittance from models of international trade— is a small but significant culprit in failures of PPP.

The duration of the paper goes as follows. First, I present relevant literature on the relationship between trade costs and PPP. Second, I explain the basic setup of my model, which draws on several of the key assumptions from models presented in articles included in the literature review; I employ various empirical techniques to support the realism of these assumptions. Finally, I build the algorithm and present conclusions.

2: Literature Review

The study of international trade has come a long way since Ricardo. Economists now realize that trade costs and frictions sometimes prevent real-world trading patterns from matching the simplistic Ricardian model of Comparative Advantage. Instead, economists recognize that if they are able to pin down each individual cost wrapped up in the term "trade costs" —there are many differentiated costs, incurred by a variety of actors, included in this term— then they will have come a long way in being able to explain international trade flows (Hummels 1999; Melitz 2003; Bernard, Jensen, and Schott 2006). As such, there is an ongoing debate over the "right" way to conceptualize trade costs.

The consensus among economists has evolved over time. In his seminal paper, Krugman (1980) proposes a model of monopolistic competition among exporting firms; he adds an Iceberg Cost component to the model, which is a feature first proposed by Samuelson (1954). Because quantities of goods deteriorate over time in this model, exporting firms must produce a greater quantity of their goods than they wish to sell in the foreign market; this effect is then more extreme over longer distances.

Could Iceberg Costs be the source of failure of PPP (or at least *a* source of failure of PPP)? Ricci and MacDonald (2002) argue that it is. This is a likely enough outcome that I include a quasi-Melting Effect in my paper's model. However, I try to add empirical rigor to the assumptions that Ricci and MacDonald (2002) make by considering the Melting Effect as a matter of inventory depreciation. Specifically, I believe the amount of value that melts away while a good is being exported is reflected in the rate of depreciation for that good; if, for example, the rate of depreciation of a certain good is 0.5% per day, I can model the value of that good as having lost 0.5% for every day it takes to transport that good to its destination. Inventory depreciation data (by industry) are published by the Bureau of Economic Analysis on a yearly basis, and thus can be analyzed empirically with relative ease.

Although theoretically tractable, the Melting Effect narrative does not hold up in the data. Anderson and Wincoop (2004) demonstrate empirically that overall trade costs comprise only a

9% equivalent tax-incidence due to the time value of goods —corresponding to the Melting Effect— compared to a 12% incidence due to international freight/transport costs, a 44% incidence due to trade barriers (tariffs and quotas), and a 55% incidence due to domestic transport and distribution costs (incurred after the goods cross the border). In their later work, Irarrazabal, Moxnes, and Opromolla (2015) corroborate these findings. However, trade barriers tend to be much less substantial in wealthy countries, for which tariffs only explain 5% of trade costs and non-tariff trade barriers only explain a further 8%. On the other hand, international shipping and domestic distribution costs tend to be a much more significant component of trade costs among wealthy countries, because profit-maximizing logistics firms charge higher price markups for these highly-demanded routes. My paper builds an empirical model of global trade based on wealthy countries only (such countries are responsible for the vast majority of global trading volumes) and thus the role of tariffs and quotas should not be as large of a consideration as otherwise. Transport costs, however, should be a key consideration for two reasons. First, as Anderson and Wincoop (2004) as well as Irarrazabal et. al (2015) show, these costs are more important in relation to trade barriers among wealthy countries; between international shipping fees and distribution costs, transport costs account for the vast majority of trade costs. Second, many papers on PPP do not consider these transport costs at all, and the few that do take shipping rates as exogenous to the model. But in reality, these transport rates —both international and domestic distribution fees— are set by profit-maximizing, monopolistically competitive shipping lines who charge a markup on their services in the same way that any other monopolistically competitive firm would.

Contemporary consensus among economists such as Bosker and Buringh (2017) and Fingleton (2007) is that overall transport costs are concave with respect to distance, and so this is another feature I try to capture. Fujita and Krugman (1995) and Krugman (1995) pioneered the idea of concave transport costs by redefining the Melting Effect from a linear decay pattern over time to an exponential decay pattern over time.

A good example of a PPP failure model with incomplete trade cost considerations is Atkeson and Burstein (2008). They attribute systematic failures of PPP to pricing-to-market behavior of exporting firms. This behavior, they argue, is especially pervasive among larger firms who have greater control over their prices. I do not see anything immediately wrong with this conclusion. But once again, the only trade cost present in this model is a Krugman-esque Iceberg Cost, which Anderson and Wincoop (2004) show to have a relatively small impact on international trade flows. And even there, Atkeson and Burstein (2008) leave this Iceberg Cost as exogenous, without recognizing that exporting firms collectively may actually have some control over the prices that shipping lines set, and vice versa.

There is also the issue of non-tradable goods, which may also cause price level inconsistencies among different countries if included in a comparative basket of goods, and thus further disruptions to PPP. This argument has a lot in common with Taylor and Taylor (2004), who remind us that even most tradable goods have at least one non-tradable input that could be distorting the internationally-comparable price of the finished product. But I view trade costs to be relevant in each of these last two arguments, for non-tradable goods are essentially goods with infinite trade costs. The algorithmic model in this paper includes several representative non-tradable goods industries for that reason.

Some have also argued for alternative causes of PPP failure that do not involve trade costs. In particular, sticky prices contribute to price level inconsistencies between countries when a “border effect” alters prices in one country to a greater extent than the other (Parsley and Wei 2001). Additionally, financial economists such as Frydman, Goldberg, Johansen, and Juselius (2009) reconcile the Purchasing Power Parity Puzzle with models of imperfect information used to understand bizarre behavior in other financial markets. Finally, some have attributed PPP failure to noise traders in foreign exchange markets, who may be distorting nominal exchange rates away from the relative price levels to which they “should” be converging (Xu 2010). There is a long list of alternative explanations along these lines that have been given. Although such alternative explanations could hold plenty of value, they are outside the scope of this paper.

3: Explanation of Model Components

The profit-maximizing decision of exporting firms in my model’s framework is based on a refined rendition of the Dixit-Stiglitz model of monopolistic competition (Dixit and Stiglitz 1977), as presented by Krugman (1980) and later Lu (2018b). In this framework, each firm produces a slightly differentiated variety z, which means firms are monopolistically competitive. Consumers in country c maximize utility according to

(1) $$U = \left(\sum_z q(z)^{\frac{\sigma_z - 1}{\sigma_z}}\right)^{\frac{\sigma_z}{\sigma_z - 1}}$$

where σ_z represents the elasticity of demand for good z. Taking the first order condition to maximize utility, subject to budget constraint Y, results in a demand function for good z of

(2) $$q(z) = \frac{{p_{z,c}}^{-\sigma_z}}{{\Pi_c}^{1-\sigma_z}} Y_c$$

where $p_{z,c}$ is the price of good z in c, Π_c is a price index for all varieties available in c, and Y_c is total consumer expenditures in c. The derivation of (2) from (1) can be found in Lu (2018a).

3.1: Profit-Maximizing Behavior of Exporting Producers

I now refine Krugman's model to incorporate the two different forms of shipping costs that firms incur: the time-sensitive Melting Effect (i.e., inventory losing some proportion of its value while enroute) present in the New Economic Geography models (see Fujita and Krugman 1995; Krugman 1995) and a distance-sensitive shipping cost.

I first consider the Melting Effect. The only difference between my Melting Effect and Krugman's is that whereas Krugman considers melting as a deterioration of *quantity*, I consider it as a deterioration of *quality*, and hence, a deterioration of *price*. This assumption is more accurate given I am using real-world depreciation rates to estimate the effect of melting; the deterioration of a good's quality over time does not decrease the *quantity* of the good, but rather it decreases the *price* that consumers are willing to pay for that quantity. For example, if I ship five apples to China, and those apples lose 50% of their value during the journey because they are not as fresh anymore, then I still have five apples. Consumers just will not be willing to pay as much for each one.

For convenience, I will omit the subscript z from here on. It takes a period of d days to ship goods to country c, so I can represent the price of a good z in country c on day t like this:

(3) $$p_{c,t} = p_{c,t-d}(1-v_t)^d\,\varepsilon_{c,t}$$

where v_t represents the proportional value that a unit of z loses on a daily basis, $p_{c,t\text{-}d}$ is the price consumers in c *would have* paid for a unit of z on day $t-d$ (the day the good is produced), and $\varepsilon_{c,t}$ represents the exchange rate of c's currency into the producer of z's home currency (producers want to be paid in their own currency). The variable v_t also represents the quality of good z, in the same spirit as Amiti, Itskhoki, and Konings (2014), where higher-quality variants have a lower v_t. Substituting (3) into (2) yields

(4) $$q_{c,t-d} = \frac{\left(p_{c,t-d}(1-v_t)^d\,\varepsilon_{c,t}\right)^{-\sigma}}{\Pi_{c,t}^{\;1-\sigma}}Y_{c,t}$$

on day $t-d$. The profit-maximizing condition reveals that monopolistically competitive firms in this model should charge a price equal to

(5)[2] $$p_{c_2,t} = \tfrac{\sigma}{\sigma-1}(1-v_t)^d\,\tfrac{1}{K_{c_2,t}}\left(m_{c_1,t} + \tau_{c_1c_2,t}\right)$$

when producing in country c_1 and selling in country c_2, where $\tau_{c_1c_2,t}$ is the per-unit transport cost to c_2 from c_1 in time t, $m_{c_1,t}$ is the unit cost of production in c_1 at time t (comprising the unit cost

[2] See Appendix A for the derivation of (5).

of labor, capital, and intermediate goods), and $K_{c_2,t}$ is a demand curve shifter equal to $\left(\frac{\Pi_{c_2,t}{}^{1-\sigma}}{Y_{c_2,t}}\varepsilon_{c_2,t}{}^{\sigma}\right)^{1-\sigma}$. Because $\Pi_{c_2,t}$, $Y_{c_2,t}$, $\varepsilon_{c_2,t}$, and σ are all assumed exogenous —individual firms do not have any significant control over these— I am able to treat $K_{c_2,t}$ as a constant when taking the first order condition. Transport cost $\tau_{c_1c_2,t}$ comprises two components: a per-unit international shipping cost $S_{c_1c_2,t}$ into country c_2 from c_1 in time *t*, and a per-unit domestic distribution cost $D_{c_2,t}$ within country c_2 at time *t*, so that $\tau_{c_1c_2,t} = S_{c_1c_2,t} + D_{c_2,t}$. Note that (5) generalizes to firms' domestic sales, for which $c_1 = c_2$ and $S_{c_1c_2,t} = 0$ for all *t*.

3.2: Reality Check: Are My Assumptions of Exporting Firm Behavior Realistic?

I want to test the reliability of equation (5) empirically, as it forms the foundation of the entire model. To do this, I seek to estimate a fixed-effects regression (with exporting country, industry, and/or time fixed-effects) to assess the quality of this assumption. (5) can be transformed to a fixed-effects regression like this:

$$E\left(p_{c_2,t,z}\right) = \beta_1 I_1 + \beta_2 I_2 + \alpha_{c_1} + \gamma_t + \delta_z \tag{5a}$$

where I_1 is an interaction term between $\frac{\sigma_z}{\sigma_z - 1}$, $\left(1 - v_{t,z}\right)^d$, $\frac{1}{K_{c_2,t}}$, and $m_{c_1,t,z}$; and I_2 is an interaction term between $\frac{\sigma_z}{\sigma_z - 1}$, $\left(1 - v_{t,z}\right)^d$, $\frac{1}{K_{c_2,t}}$, and $\tau_{c_1c_2,t}$. (5) can also be rewritten as:

$$p_{c_2,t,z} = \beta_1 \frac{\sigma_z}{\sigma_z - 1}\left(1 - v_{t,z}\right)^d \frac{1}{K_{c_2,t}} m_{c_1,t,z} + \beta_2 \frac{\sigma_z}{\sigma_z - 1}\left(1 - v_{t,z}\right)^d \frac{1}{K_{c_2,t}} \tau_{c_1c_2,t} + \alpha_{c_1} + \gamma_t + \delta_z \tag{5b}$$

I then seek to test whether β_1 and β_2 are statistically significant, as well as whether the inclusion of I_1 and I_2 in the regression causes a significant increase in R^2 over a regression that includes all relevant variables, but excludes the interaction terms.

First, a few notes on my empirical methodologies. The data include monthly goods sales in the United States on goods from seven different markets of productive origin (China, Japan, United Kingdom, Euro Area, Canada, Mexico, and the United States). Import good price indices, which constitute the dependent variable, come from FRED. There are 47 country-industry pairs. An additional country fixed effect is not needed for c_2 because my dataset only includes imports into the US, and thus c_2 is constant. Hence, the only country fixed effect included is for the exporting country, c_1. $K_{c_2,t}$ is calculated using FRED time-series data on consumer expenditures (representing $Y_{c_2,t}$), CPI (representing $\Pi_{c_2,t}$), and nominal exchange rates (representing $\varepsilon_{c_2,t}$); estimates of industry-specific σ_z come from Miklosovic and Lichner (2011), Alaouze (1977), Akinci (2017), and Salem (2005). Industry- and year-specific depreciation rates on consumer durable goods (which represent $v_{t,z}$) come from Bureau of Economic Analysis figures; estimates of d, or the number of days spent enroute to the destination market, come from the Freightos Transit Time Calculator. I estimate $m_{c_1,t,z}$ empirically in the same manner as Hall (2018)[3], using data from the KLEMS World database[4]. Most critically, I estimate $\tau_{c_1 c_2,t} = S_{c_1 c_2,t} + D_{c_2,t}$ in the following manner. For the international $(S_{c_1 c_2,t})$ component, I use a weighted average of route-specific international maritime freight rates from the Freightos Baltic Index (FBX)[5], route-specific international air freight rates from Freight Analysis Framework, and, for applicable countries[6] (Mexico and Canada), route-specific international trucking, rail, and pipeline freight rates from the Bureau of Transportation Statistics, Freight Analysis Framework, and CEIC. Tariffs, duties, dock handling fees, and other surcharges are baked into the FBX, and are not explicitly defined in the regression for that reason. Data on relative importances of these five modes of transport, which allow me to weight the prices of each mode accurately, come from the US Bureau of Transportation Statistics and Freight Analysis Framework. As for the domestic $(D_{c_2,t})$ component, I use a weighted average of long-distance truckload (TL) trucking[7], rail, pipeline, air, and vessel freight rates within the US from FRED and Freight Analysis Framework. Finally, I weight $S_{c_1 c_2,t}$ and $D_{c_2,t}$ to be consistent with their relative importances according to Anderson and Wincoop (2004).

It appears that this model is robust to endogeneity. Simultaneity is avoided as the price that a single monopolistically-competitive firm chooses is not a significant *predictor* of depreciation, macroeconomic conditions in the destination market, production costs, or transport costs, but rather a *response* to each of these; omitted variable bias is also largely avoided as the model accounts for the vast majority of significant forms of trade costs that firms in wealthy countries face, according to Anderson and Wincoop (2004). To confirm this assumption, I run an exogeneity test on a hypothetical first-stage regression similar to (5a) —the only difference being

---

[3] Marginal Costs are estimated as the ratio of the change in total cost to the change in output.

[4] I assume that marginal costs for a given industry in a given country are time-invariant apart from the change caused by shifts in exchange rates, because of time inconsistencies between the KLEMS data and the FBX data. This assumption is a little simplistic. However, due to these time-period inconsistencies I would either have to assume that transport costs are time-invariant —which would completely defeat the purpose of this paper— or assume that marginal costs are quasi-time-invariant, in order to estimate my regressions.

[5] Freightos provides weekly container shipping rate data, disaggregated by shipping route (e.g., Northern Europe to North America, East Asia to Northern Europe, etc.), for trade among East Asia, North America-East Coast, North America-West Coast, and North Europe going back to 1Q2018; international shipping rates on imports from these areas from before that are only reported annually and originate from Freight Analysis Framework.

[6] Maritime and air freight accounts for all imports from European and Asian countries to the US, so I do not have to factor in other modes for these countries. Overall, maritime freight accounts for 90% of international volumes.

[7] I choose to exclude less-than-truckload (LTL) trucking rates for two reasons: first, LTL rates are highly regulated by the National Motor Freight Traffic Association (NMFTA) and are thus less likely to adjust predictably to open-market pricing determinants; second, TL covers higher-volume movements, and is thus a better proxy for the trucking rates that I am trying to capture in my model (Spelic 2017).

the lack of a time fixed effect[8]— using the Stata command **estat endog**. This test results in a Wu-Hausman statistic of 0.178 and a Durbin statistic of 0.181, both well within the $p \geq 0.05$ significance threshold.

If (5) is a realistic assumption, we should see significant coefficients on interaction terms I_1 and I_2, and ideally though not necessarily, R^2 should increase in comparison with a model that excludes these interaction terms. In other words, I do not simply want to prove that there is *some* effect of $\frac{\sigma_z}{\sigma_z - 1}$, $\left(1 - v_{t,z}\right)^d$, $\frac{1}{K_{c_2,t}}$, $m_{c_1,t,z}$, and $\tau_{c_1 c_2,t}$ on the prices that exporting firms set, but that each of these predictors has *the particular* interactive effect on prices that (5) describes. The magnitude of the coefficients does not matter as much because the dependent variable is made up of price *indices*, not the prices themselves. Tables 3.1 and 3.2 show empirical support for (5).

Table 3.1: Empirical Support for (5)

Regressor	**Dependent Variable:** $p_{c_2,t,z}$							
$\frac{\sigma_z}{\sigma_z - 1}$	0.0039 (0.0038)		-0.1066*** (0.0352)	1.012*** (0.2718)		-1.6389*** (0.2488)	-0.0103*** (0.0036)	
$\left(1 - v_{t,z}\right)^d$	2664.35*** (575.961)		3080.26*** (592.769)	148737*** (22471.8)		196023*** (20406.7)	2440.60*** (534.501)	
$\frac{1}{K_{c_2,t}}$	-0.0085 (0.0010)		-0.0373*** (0.0137)	0.3065*** (0.0257)		0.1662*** (0.1685)	0.0032 (0.0099)	
$m_{c_1,t,z}$	-0.0005*** (4.04E-5)		-0.0006*** (4.57E-6)	0.0013*** (0.0001)		0.0023*** (0.0001)	-0.0005*** (3.91E-5)	
$\tau_{c_1 c_2,t}$	0.1330*** (0.0452)		0.1067*** (0.0454)	0.0835** (0.0370)		0.0302 (0.0335)	-0.0282 (0.0447)	
I_1		-2.6E-6** (1.16E-6)	5.14E-6*** (1.26E-6)		1.18E-5*** (8.70E-6)	2.14E-5*** (1.11E-6)		-2.65E-6** (1.08E-6)
I_2		3.51E-5 (3.65E-5)	0.0009*** (0.0003)		0.0036*** (0.0003)	0.0039*** (0.0004)		-7.74E-5** (3.45E-5)
Industry FE	No	No	No	Yes	Yes	Yes	No	No
Country FE	No	No	No	No	No	No	Yes	Yes
Time FE	No	No	No	No	No	No	No	No
Adj. R^2	0.0889	0.0160	0.0995	0.4317	0.4433	0.5373	0.2516	0.1989
Observations	1,953	1,953	1,953	1,953	1,953	1,953	1,953	1,953

*, **, and *** represent significance on a 90%, 95%, and 99% confidence interval, respectively

[8] I believe this regression (represented by the fourth column from the left in Table 3.2) to be the best assessment of (5) out of any of the regressions that I run.

Table 3.2: Empirical Support for (5) (cont.)

Regressor	**Dependent Variable:** $p_{c_2,t,z}$							
$\frac{\sigma_z}{\sigma_z - 1}$	-0.0301 (0.0386)	-0.0906*** (0.0348)	-1.3548*** (0.2561)		-1.7314*** (0.2398)	-0.5626** (0.2686)		-1.0320*** (0.2529)
$\left(1 - v_{t,z}\right)^d$	2624.05*** (568.956)	3009.89*** (584.769)	154406*** (20933.4)		188671*** (19493.9)	84919.8*** (22120.8)		128223*** (20717.7)
$\frac{1}{K_{c_2,t}}$	0.0007 (0.0129)	-0.0333*** (0.0136)	0.0658** (0.0292)		0.1197*** (0.0346)	0.0576** (0.0287)		0.1216*** (0.0341)
$m_{c_1,t,z}$	-0.0005*** (4.68E-5)	-0.0006*** (4.55E-5)	0.0033*** (0.0002)		0.0041*** (0.0002)	0.0030*** (0.0002)		0.0038*** (0.0002)
$\tau_{c_1 c_2,t}$	-0.0292 (0.0445)	0.1655*** (0.0606)	0.0295 (0.0362)		0.0292 (0.0336)	-0.0686 (0.0618)		-0.0852 (0.0574)
I_1	6.18E-6*** (1.28E-6)	5.05E-6*** (1.26E-6)		0.0001*** (1.18E-6)	1.93E-5*** (1.08E-6)		1.19E-5*** (1.14E-6)	1.84E-5*** (1.07E-6)
I_2	0.0001 (0.0003)	0.0008*** (0.0018)		0.0024*** (0.0004)	0.0018*** (0.0005)		0.0022*** (0.0004)	0.0014** (0.0005)
Industry FE	No	No	Yes	Yes	Yes	Yes	Yes	Yes
Country FE	Yes	No	Yes	Yes	Yes	Yes	Yes	Yes
Time FE	No	Yes	No	No	No	Yes	Yes	Yes
Adj. R^2	0.2598	0.1246	0.5174	0.4621	0.5863	0.5338	0.4978	0.5986
Observations	1,953	1,953	1,953	1,953	1,953	1,953	1,953	1,953

*, **, and *** represent significance on a 90%, 95%, and 99% confidence interval, respectively

Overall, (5) does not seem too far off base from an empirical standpoint. All three fixed effects seem to have a positive impact on my ability to explain price markups on exported goods, although the time fixed effect is a little less meaningful than the other two. Most importantly, the interaction terms are almost always significantly positive, and their inclusion almost always causes a significant increase in R^2 over the comparable model that excludes them; the interaction terms cause an especially significant increase in R^2 of over 0.1 when an industry fixed effect is introduced to the model.

There are of course some limitations to this empirical study. First and foremost, I use an unbalanced panel because of inconsistencies in transport cost data. Specifically, the FBX only goes back to early 2018, and thus data for goods from China, Japan, the Eurozone, and the UK only go back to early 2018; however, data for goods from Mexico, Canada, and the US go back to 2012 (I was unable to find a reliable data source for route-specific international maritime freight rates from before 2018). The unbalanced panel may bias my results somewhat. However, I feel it is still preferable to deleting all records from before 2018, as I only had around 2,000 observations to begin with. Second, as will be explored further in section 4, the time period of this empirical study is different from the time period of the algorithm I eventually build. This discrepancy is due to time limitations on import price indices in the US, which only go back to June 2012. It is a possibility that empirical evidence supporting (5) would not be as strong for 1999-2008 (the period on which I construct the algorithmic model) as it is for 2012-2019. Additionally, time limitations on the availability of KLEMS data may confound the estimated effect of $m_{c_1,t,z}$ and/or I_1; this may explain why the coefficient on $m_{c_1,t,z}$ is sometimes negative, which is a highly counterintuitive result. Third, there are only reliable data on imports into the US; I only study US imports here as a result of that, but a more thorough examination of (5) would incorporate other countries' imports as well, perhaps with an additional fixed effect for importing country. And lastly, four-way interaction terms are obviously very unusual and do not give me any significant insight as to the individual importance of each of the factors affecting export price markups. The only purpose the interaction terms serve here is to grant empirical support to (5), which I have more or less done.

This back-of-the-envelope empirical study does not constitute the central finding of this paper. Rather, it is a high-level assessment of whether the assumption presented by (5), knowing that it is tractable, is also at least relatively realistic.

3.3: To Export Or Not To Export?

A firm only exports to a certain country on a certain day if it thinks it can make a profit by doing so. But this perception is based on the firm's expectations of market conditions when the good arrives in the destination country, *not* what they are now (i.e., quantity demanded $q_{c_2,t}$ may not be perfectly predictable on day $t-d$). In other words, although firms are ultimately price-makers of $p_{c,t}$, they cannot adjust to current market conditions the *quantity* they export to a given country after the shipment has already departed. This is a concept borrowed from Duarte and Stockman (2001). Hence, a firm will only ship to country c_2 on day $t-d$ if:

$$(6) \qquad Pred\left(p_{c_2,t}q_{c_2,t}\right)-\left(m_{c_1,t-d}+\tau_{c_1c_2,t-d}\right)Pred\left(q_{c_2,t}\right)>0$$

where $m_{c_1,t-d}$ is the unit cost of production on day $t-d$ in c_1. $Pred\left(p_{c_2,t}q_{c_2,t}\right)$ represents a firm's predicted revenue (price multiplied by quantity) from selling in c_2, d days in the future; $Pred\left(q_{c_2,t}\right)$ represents the quantity that a firm produces on day $t-d$ because they want to produce a quantity equal to the amount they expect to be demanded in c_2 on day t. This relationship originates from the assumption that if a firm incurs its costs of production by producing today, it will not enjoy any revenue from selling in c_2 what they produce today in c_1

until d days from now. Written out fully by substituting the profit-maximizing $p_{c,t}$ and $q_{c,t}$ into (6) and reducing predicted values, a firm will only ship to country c_2 on day $t-d$ if:

$$(7)^{9} \quad \left(m_{c_1,t-d}+\tau_{c_1c_2,t-d}\right)^{1-\sigma}\left(\left(\frac{\sigma}{\sigma-1}\right)^{1-\sigma}(1-v_{t-d})^{d-d\sigma}Pred\left(K_{c_2,t}{}^{\sigma-1-\frac{1}{1-\sigma}}\right)-\left(\frac{\sigma}{\sigma-1}(1-v_t)^{d}\right)^{-\sigma}Pred\left(K_{c_2,t}{}^{\sigma-\frac{1}{1-\sigma}}\right)\right)>0\,.$$

How accurately are producers able to make this prediction? Because σ is constant across time, the firm can always predict σ perfectly ahead of time, no matter how far into the future they are trying to predict. For that reason, the quality of the producers' forecasts depends completely on their ability to predict $K_{c_2,t}$, which represents a demand-curve-shifting parameter.

Empirical evidence is very promising in terms of the accuracy of such predictions. For one, a simple observation of $K_{c_2,t}$ across various goods and trade routes shows that K_{c_2} tends to be very slow-moving over time in the vast majority of cases. This would make intuitive sense, as the macroeconomic variables that $K_{c_2,t}$ embeds tend to be relatively stable for the wealthy countries that this paper studies. And recall that σ —which is assumed time-invariant in this model— is also one of the determinants of $K_{c_2,t}$. As a result, it would seem *a priori* that $K_{c_2,t-d}$ is a strong predictor of $K_{c_2,t}$. But I want to add a little more rigor to this observation. To do so, I calculate the autocorrelation of $K_{c_2,t}$ (which producers try to predict) against $K_{c_2,t-d}$ (which producers know) for all 1,582 route-good combinations[10] for which I have data. This autocorrelation is extremely strong across the vast majority of route-good combinations. Figure 3.1 shows the distribution of lag-*d* autocorrelations in the sample.

Figure 3.1: Distribution of Autocorrelations in K

Lag d

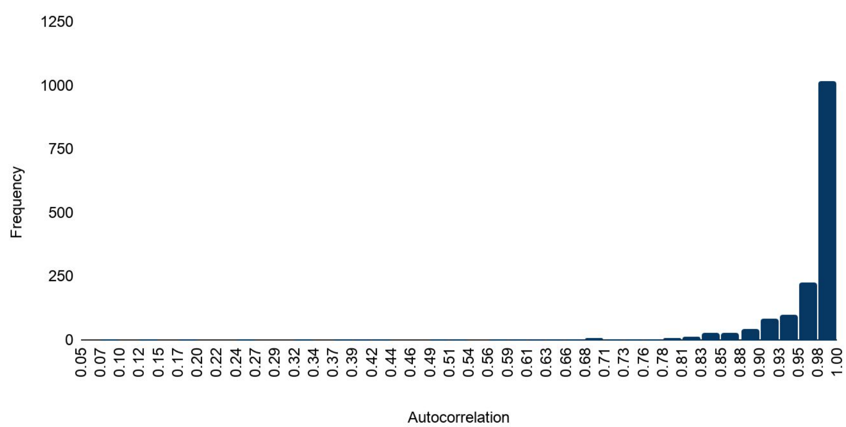

[9] See Appendix B for the derivation of (7).

[10] A route-good combination signifies a certain good being shipped from a certain origin to a certain destination. For example, fruits and vegetables being shipped from China to Canada is a route-good combination; however, this is distinct from fruits and vegetables being shipped from Canada to China, which is a separate route-good combination. Each route-good combination has a unique K_t.

Furthermore, the autocorrelation remains quite strong even when trying to predict farther into the future than d days from now. Figure 3.2 shows the distribution of lag-2d autocorrelations.

Figure 3.2: Distribution of Autocorrelations in K

Lag 2d

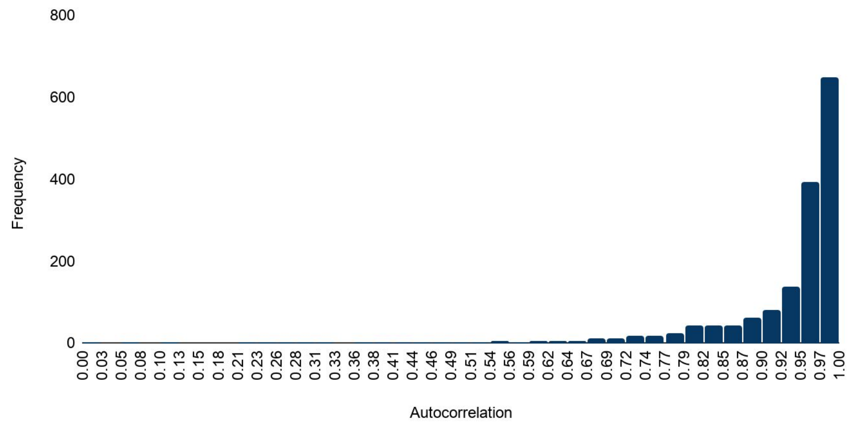

Expectedly, autocorrelations tend to drop slightly as one tries to predict farther off in the future. However, there is still a very strong autocorrelation in this case, even multiple periods in the future. These two figures show us that K_{c_2} displays strongly non-stationary behavior; in other words, not only does $K_{c_2,t-d}$ tend to be a nearly perfect predictor of $K_{c_2,t}$, but so do $K_{c_2,t-2d}$, $K_{c_2,t-3d}$, and so on for a good long while. As a result, each of these known values provides very strong predictive value to firms as to what conditions will be in destination markets when their shipments arrive.

The point I am getting at here is that firms can predict $K_{c_2,t}$ with near-perfect accuracy. A period of d days is generally not enough time for K_{c_2} to change significantly among stable, wealthy economies. Hence, I will assume going forward that firms can make these predictions perfectly. It simply does not take enough time for goods to be shipped to destination markets, that any significant inefficiencies in trade patterns could be caused by uncertainty of future demand conditions. In the model, this assumption shows up as $Pred\left(p_{c_2,t}q_{c_2,t}\right)=p_{c_2,t}q_{c_2,t}$ and $Pred\left(q_{c_2,t}\right)=q_{c_2,t}$, which means (7) can be rewritten as:

$$(8)\quad \left(m_{c_1,t-d}+\tau_{c_1c_2,t-d}\right)^{1-\sigma}\left(\left(\frac{\sigma}{\sigma-1}\right)^{1-\sigma}(1-v_{t-d})^{d-d\sigma}K_{c_2,t}^{\ \sigma-1-\frac{1}{1-\sigma}}-\left(\frac{\sigma}{\sigma-1}(1-v_t)^{d}\right)^{-\sigma}K_{c_2,t}^{\ \sigma-\frac{1}{1-\sigma}}\right)>0$$

.

3.4: Profit-Maximizing Shipping Lines

Exporting firms take $\tau_{c,t\text{-}d}$ for each c as a given. $\tau_{c,t\text{-}d}$ is the price determined by a profit-maximizing representative shipping firm who is responsible for shipping goods to their destination markets. I assume that one representative firm handles both international shipping and domestic distribution, as trying to model each part separately would be mathematically intractable. Suppose that these shipping lines know the general interests of exporting firms, but cannot predict whether each firm will want to export any goods on day $t-d$. On day $t-d$, producers will demand the following quantity of shipping from country c_1 to country c_2:

(9)[11]
$$Q_{c_1c_2,t-d} = \sum_{\zeta} \frac{p_{\zeta,t-d}^{\;-\sigma_\zeta}}{K_{c_2,t-d}} = \sum_{\zeta} \frac{\left(\frac{\sigma_\zeta}{\sigma_\zeta-1}\left(\frac{1}{K_{c_2,t-d}}\right)\left(m_{c_1,t-d,\zeta}+\tau_{c_1c_2,t-d}\right)\right)^{-\sigma_\zeta}}{K_{c_2,t-d}}$$

where each ζ represents a good whose producer in c_1 is interested in exporting to c_2 on day $t-d$ (the set of all ζ's is a subset of all z's). The profit of the shipping line that results from shipping a total quantity of $Q = \sum_{\zeta} q_\zeta$ goods between c_1 and c_2 departing at time $t-d$ is then:

(10)
$$\pi_{c_1c_2,t-d} = \tau_{c_1c_2,t-d}Q_{c_1c_2,t-d} - \sum_{\zeta}\theta_{c_1c_2,t,\zeta}q_{\zeta,t-d} =$$
$$\tau_{c_1c_2,t-d}\sum_{\zeta}\frac{\left(\frac{\sigma_\zeta}{\sigma_\zeta-1}\left(\frac{1}{K_{c_2,t-d}}\right)\left(m_{c_1,t-d,\zeta}+\tau_{c_1c_2,t-d}\right)\right)^{-\sigma_\zeta}}{K_{c_2,t-d}} - \sum_{\zeta}\frac{\theta_{c_1c_2,t,\zeta}\left(\frac{\sigma_\zeta}{\sigma_\zeta-1}\left(\frac{1}{K_{c_2,t-d}}\right)\left(m_{c_1,t-d,\zeta}+\tau_{c_1c_2,t-d}\right)\right)^{-\sigma_\zeta}}{K_{c_2,t-d}} =$$
$$\sum_{\zeta}\left(\tau_{c_1c_2,t-d} - \theta_{c_1c_2,t,\zeta}\right)\frac{\left(\frac{\sigma_\zeta}{\sigma_\zeta-1}\left(\frac{1}{K_{c_2,t-d}}\right)\left(m_{c_1,t-d,\zeta}+\tau_{c_1c_2,t-d}\right)\right)^{-\sigma_\zeta}}{K_{c_2,t-d}}$$

where θ is a concave function of d that describes the cost of physically transporting one unit of ζ from c_1 to c_2 (the function is concave to be consistent with Bosker and Buringh (2017), Fingleton (2007), Fujita and Krugman (1995), and Krugman (1995)). This presents a key distinction in the model: the only transport cost that goods-producing firms incur is τ, which represents the financial expense of paying a shipping line to transport goods; meanwhile, the shipping lines only incur θ (representing the economic cost of physically transporting each good) which may be very distinct from τ.

Notice also that θ depends on ζ, but that τ does not. This is consistent with real-world patterns. In the real world, the price of a 40-foot intermodal shipping container being sent from

[11] (9) is based on a summation of (2), which evaluates the quantity exported along each trade route by each individual firm.

some origin to some destination is usually standardized, regardless of its contents (Redwood Logistics 2019). Thus, although the shipping line transports many varieties of goods with differing θ's, the shipping line in this model must select a single price τ for the shipment of all goods from c_1 to c_2 on day $t-d$.

The objective of the shipping line is to set their price to $argmax_{\tau_{c_1c_2,t-d}}(\pi_{c_1c_2,t-d})$, or the price that will maximize profits of shipping from c_1 to c_2. But shipping lines cannot predict which producers will be trying to ship on that day, given that producers make their exporting decisions *in response to*, and not *ahead of*, the prices that shipping lines set. Hence, the shipping line assumes that the producer of every good z wants to ship on that day. Taking the first order condition of (9) reveals that there is no closed-form solution[12] for $argmax_{\tau_{c_1c_2,t-d}}(\pi_{c_1c_2,t-d})$ —or, at least, not one that can be found using any math with which I am familiar— and so in this model, I employ a Gradient Ascent Search algorithm to calculate the "optimal" price (gradient ascent search is not always optimal, but is the best option available to me given what I am trying to model). However, if at the optimal shipping price profits for the transport firm for a particular route are still negative, the transport firm will not make a shipment along that route on day $t-d$, and the quantity of all goods arriving in country c_2 from country c_1 on day t will be 0. Hence, in order for *any* good z to be shipped from c_1 to c_2 on day $t-d$ and arrive on day t, two things must be true:

(a) It must be profitable for z's producer to make such a shipment, i.e., (8) must hold, AND
(b) It must be profitable for the shipping line to send a shipment from country c_1 to country c_2 on day $t-d$, i.e., it must hold that $\pi_{c_1c_2,t-d}\left(argmax_{\tau_{c_1c_2,t-d}}\right) > 0$.

What is the best way to model the function θ? Unlike the assumptions presented by (5) and (7), I do not have the necessary data to ensure the realism of any assumption I make of θ. It would be insufficient to consider determinants of the FBX, because the FBX only measures shipping *prices*, not measure specifically the *cost of physically transporting the goods*; i.e., the FBX measures τ which, as aforementioned, may be very distinct from θ. Instead, I am forced to model θ using intuition only. I use:

(11) $$\theta_{c_1c_2,t,\zeta} = \delta_\zeta \omega_t \sqrt{d}$$

where δ_ζ is a time-constant scalar that represents the difficulty of transporting a unit of ζ, and ω_t represents the price of oil at time t. To reiterate, (11) is consistent with θ being a concave function of d. Because I do not have a good way to estimate δ_ζ empirically, I will assume for the purposes of the simulation that $\delta_\zeta \sim N(5, 1)$[13].

---

[12] See Appendix C

[13] I choose 5 as the median/mean because a simple test of (11) with $\delta_\zeta = 5$ and oil prices consistent to what they were on 3/1/20 produced τ_c roughly equal to real-world shipping prices on that day.

3.5: Domestic Price Levels

To figure out price levels in each country, it is insufficient to simply calculate the price of the overall basket of goods available in that country at that time. This would be represented by

(12) $$P_{c,t} = \sum_{z} q_{z,t} p_{z,t} \,.$$

where z is a good available for purchase in country c. This is insufficient because more varieties —and higher quantities of those varieties— will be available in countries with higher Y_c, and thus a direct comparison of countries with varying Y_c using (12) would be biased. Instead, I estimate price level in country c_2 at time t as the ideal price index demonstrated by Lu (2018b):

(13) $$P_{c_2,t} = \sum_{z} \left(\sum_{c} q_{c,t,z} p_{c,t,z}{}^{1-\sigma_z} \right)^{\frac{1}{1-\sigma_z}}$$

where $q_{c,t,z}$ and $p_{c,t,z}$ represent the quantity and price respectively of good z exported to c_2 from every individual country c (sum over all countries exporting to c_2, including c_2 itself). This is somewhat of an abuse of notation, so let me clarify: in this case, $q_{c,t,z}$ and $p_{c,t,z}$ represent the quantity and price respectively of good *z in country* c_2 among units of z produced in *specifically in c*; $q_{c,t,z}$ and $p_{c,t,z}$ *do not* represent the overall price and quantity of good z in c at time t.

3.6: Relative Price Levels & Purchasing Power Parity Conditions

PPP conditions hold between a pair of countries $\{c_1, c_2\}$ in time t if the relative price level of that pair of countries is equal to their nominal exchange rate at that time. To figure out whether this property holds for each pair of countries in the model, I simulate price levels for each country on each day algorithmically according to (13) and compare it to the observed nominal exchange rate on that day. To make my point a little clearer, I parameterize each relative price level by a constant scalar $\lambda_{c_1 c_2}$ to ensure that for each pair of countries, PPP conditions are satisfied on the first day of the simulation. However, $\lambda_{c_1 c_2}$ does not change as time progresses. In other words, relative price level between c_1 and c_2 is calculated as:

(14) $$\frac{P_{c_2,t}}{P_{c_1,t}} = \lambda_{c_1 c_2} \frac{\sum_{z} \left(\sum_{c} q_{c,t,z} p_{c,t,z}{}^{1-\sigma_z} \right)^{\frac{1}{1-\sigma_z}}_{2}}{\sum_{z} \left(\sum_{c} q_{c,t,z} p_{c,t,z}{}^{1-\sigma_z} \right)^{\frac{1}{1-\sigma_z}}_{1}}$$

where the numerator of the right-hand side describes price level according to (13) in c_2 and the denominator describes the same thing for c_1.

4: Data and Algorithms

The algorithmic model can be thought of as a series of strategic decisions made by each player every day, starting on day 0. There are seven countries/geographies in the model: the United States, Mexico, Canada, the United Kingdom, the Eurozone[14], China, and Japan. The data on which this simulation depends range from 1999-2008[15], or 3,653 consecutive days. There are up to 35 tradable goods for each country in the model, as well as three non-tradable goods: water, electricity, and gas[16]. There are also two representative service-sector industries in each country, one tradable and one non-tradable. It is assumed that for all services, transport costs are 0. Non-tradable goods and services can only be produced domestically, but tradable goods can be produced domestically and/or internationally. Data for each variable needed to run the simulation originate from the same sources as in section 3.2. I simulate ω_t with the WTI Crude price.

On each day, the shipping line acts first. They set shipping prices between each pair of countries using (10). Because there is no closed-form solution[17] for optimal price according to (10), I find a quasi-optimal price using Gradient Ascent Search. This is a very computationally expensive operation, as the function runs in $O(pn^2)$ time for p ordered pairs of countries and n possible prices considered. I am thus held to considering no more than 60 possible prices for each route on each day, as doing so takes around twenty seconds for the algorithm to compute; this operation must be repeated on all 3,653 days, and so allowing any more time for it to compute would be impractical. From there, the shipping line can set prices for each route, and decide as to whether to send a shipment along each route that day. They face a tradeoff when choosing their prices: if they increase their prices, then they increase revenue per unit shipped, but they also discourage firms from exporting at all.

Once shipping prices for each route have been calculated, the exporting firms act. First, each one decides whether or not to export to each country according to (8); if they decide to export to said country, they determine the quantity to export and the price at which they sell their goods according to (5) and (2), respectively (recall the assumption that firms can predict quantity demanded perfectly, so the market for each good in this model always clears).

---

[14] In this model, the "Eurozone" consists of the original Eurozone countries, with the exception of Portugal: i.e., Austria, Belgium, Finland, France, Germany, Ireland, Italy, Luxembourg, the Netherlands, and Spain.

[15] I was unable to find productivity data on several of these geographies on any more recent a time frame than 2008-2009, which is why it can go no later; the Euro regime began in 1999 which is why it can go no earlier.

[16] Countries have the following number of representative tradable goods: United States 35, China 35, Japan 34, Eurozone 35, United Kingdom 33, Canada 34, Mexico 29. Certain goods are excluded in certain countries when productivity data is unavailable.

[17] See Appendix C

I can then calculate price level in each country according to (13), but only using goods that have already arrived in their destination markets (and thus are available for consumption). Finally, I compare these simulated price levels between countries on each day, which gives me an estimation of relative price level, which is SRPL; this number can then be compared to nominal exchange rates to determine whether or not PPP conditions hold.

4.1: Pseudocode

Written in pseudocode, the algorithm[18] for the case in which shipping lines act strategically (the second case) appears as:

```
1   FOR each day t[0,T]:
2       FOR each ordered pair of countries {c_1,c_2}:
3           s = optimal τ from c_1 to c_2 using Gradient Ascent Search, according to (10)
4           IF π(s) > 0:
5               Open route c_1c_2 to trade on day t
6           ELSE:
7               Close off route c_1c_2 to trade on day t
8           FOR each producer in c_1:
9               IF (8) holds for that producer on day t:
10                  p = price set by that producer in c_2 according to (5)
11                  q = quantity shipped to c_2 on day t according to (2)
12                  indicate that that shipment does not arrive until day t + d
13      FOR each country c:
14          G = set of all available goods that have already arrived in c by day t
15          P_c = price level in c on day t according to (13) if G are all available in c
16          Remove each g ∈ G from availability because it has been "consumed"
17      FOR each pair of countries {c_1,c_2}:
18          R = relative price level of c_1 and c_2 according to (14) with known {P_1,P_2}
19          Compare R to nominal exchange rate
```

In the first case, strategic shipping firms are omitted. The pseudocode for this case is identical to this second case, but lines 3-7 are omitted and $\tau = 0$ for all $\{c_1,c_2\}$, for all t, in equation (5) (which is implemented on line 10). Additionally, the first case assumes that all shipping routes are open every day.

[18] I implement the algorithm in Python using the Jupyter development environment.

5: Algorithm Results and Conclusions

Tables 5.1 and 5.2 show the R^2 between SRPL and NER for all 21 pairs of countries in the model, before and after I incorporate profit-maximizing shipping firms, respectively.

5.1: First Case: Melting Effect Only

Table 5.1

Country	USA	CHN	JPN	EUR	GBR	CAN	MEX
USA	N/A	0.20	0.15	0.17	0.23	0.38	0.43
CHN	0.20	N/A	0.09	0.13	0.06	0.12	0.39
JPN	0.15	0.09	N/A	0.31	0.27	0.20	0.20
EUR	0.17	0.13	0.31	N/A	0.11	0.37	0.36
GBR	0.23	0.06	0.27	0.11	N/A	0.01	0.15
CAN	0.38	0.12	0.20	0.37	0.01	N/A	0.28
MEX	0.43	0.39	0.20	0.36	0.15	0.28	N/A

Among the 21 country pairings, the median R^2 for this case is 0.20.

5.2: Second Case: Melting Effect and Strategic Shipping Lines

In Table 5.2, the R^2 values that are lower than that of the same pair of countries in Table 5.1 are highlighted in red; those that are higher are highlighted in green. Therefore, values highlighted in red correspond to pairs of countries in which profit-maximizing shipping firms make PPP satisfaction less likely, and vice versa for values highlighted in green. Among the 21 country pairings, the median R^2 for this case is 0.12, which is 0.08 lower than in the first case.

Table 5.2

Country	USA	CHN	JPN	EUR	GBR	CAN	MEX
USA	N/A	0.19	0.12	0.10	0.19	0.49	0.09
CHN	0.19	N/A	0.02	0.04	0.01	0.04	0.21
JPN	0.12	0.02	N/A	0.16	0.28	0.12	0.07
EUR	0.10	0.04	0.16	N/A	0.36	0.22	0.30
GBR	0.19	0.01	0.28	0.36	N/A	0.19	0.05
CAN	0.49	0.04	0.12	0.22	0.19	N/A	0.12
MEX	0.09	0.21	0.07	0.30	0.05	0.12	N/A

With a few exceptions, R^2 is lower in the second case than the first for the majority of country pairings: 17 out of 21 total. Interestingly, three out of the four instances in which this is not the case involve the United Kingdom; I do not know whether this is just a coincidence, or if there is actually something significant about the UK that changes the effect of for-profit shipping on price level there. And most crucially, the median difference of R^2 by 0.08 entails that generally speaking, 8% of deviation from PPP conditions can be explained by profit-maximizing behavior in shipping markets.

5.3: Marginal Effects of Shipping Profit-Maximizing on PPP Failure

The analysis in sections 5.1 and 5.2 is somewhat incomplete in that it does not capture how *changes* in shipping behavior affect the likelihood of PPP satisfaction. To address this question, I conduct a post-analysis of algorithm output data to see if there is a correlation between simulated shipping price markups and PPP likelihood. I measure shipping price markups as the logged ratio of τ to average θ for all goods that the shipping firm transports on a given day on a given route; I measure PPP likelihood as changes in mean absolute percentage error between SRPL and NER[19]. There are lots of other factors that also influence SRPL in the algorithm (e.g., consumer expenditures, demand elasticities, depreciation, etc.) and I want to weed out the effect of each of these variables on PPP likelihood so that I can focus specifically on the effect of shipping behavior. The data representing the vast majority of these features only update monthly, if at all. As a result, almost all of the within-month variation in SRPL can be

[19] e.g., if the mean absolute percentage error between SRPL and NER goes down by 5% on some day, then PPP satisfaction has become 5% more likely

attributed to changes in τ and/or θ. Hence, the regression coefficients I estimate only measure within-month variation in PPP likelihood[20].

If shipping firms were not profit-maximizers, the ratio of τ to average θ would always be 1 because a profit-non-maximizer would not set a price markup over its unit costs. Thus, shipping price markups in excess of 1 are a comprehensive evaluation of profit-maximizing behavior within shipping markets, and the correlations I study here are representative of the marginal effect of such behavior on PPP likelihood. Table 5.3 shows the regression coefficients associated with these correlations, for each ordered country pairing. Note that unlike in Tables 5.1 and 5.2, the ordering of the countries in each pair matters in Table 5.3 because neither τ nor θ is usually directionally symmetric on a given day[21]. The interpretation of the values in Table 5.3 is the number of percentage points more likely PPP satisfaction becomes for that trade flow, for every 1% that τ is increased in relation to θ. For example, a regression coefficient of 2.1319 in the box with China as the exporter and the US as the importer would signify that for that specific trade flow, PPP satisfaction becomes 2.139 percentage points less likely for every 1% increase in the shipping price markup on that route.

Table 5.3

	Exporting Country:	USA	CHN	JPN	EUR	GBR	CAN	MEX
Importing Country:								
USA		N/A	2.1319** (1.0218)	0.1368*** (0.0041)	-0.0443 (1.1187)	1.5870 (1.7270)	1.2720*** (0.4098)	1.0080*** (0.3375)
CHN		-2.1413** (0.9336)	N/A	1.5256*** (0.0387)	-0.0462 (0.0680)	-0.6638 (0.7950)	1.8325*** (0.2125)	-1.3178 (0.9415)
JPN		0.7390*** (0.2366)	-1.1652** (0.5015)	N/A	-0.4262 (0.6997)	0.9149 (1.0335)	1.1231*** (0.4695)	-0.5997 (0.5335)
EUR		0.6149*** (0.1116)	0.4415*** (0.0150)	0.0071** (0.0034)	N/A	-1.4695** (0.7402)	1.6921 (1.3625)	-0.8246*** (0.0156)
GBR		-2.0400 (1.7220)	0.8502*** (0.0507)	0.0501*** (0.0057)	1.7158** (0.7399)	N/A	1.6855*** (0.1221)	1.0284*** (0.0488)
CAN		1.6217*** (0.5635)	-2.3748 (2.1253)	-0.1213*** (0.0043)	3.4497** (1.3691)	-1.4840 (2.1301)	N/A	0.5098*** (0.0206)
MEX		-0.7185 (0.8363)	4.2200*** (0.8398)	1.3220*** (0.0284)	1.8494*** (0.5025)	0.9660 (0.8069)	-0.8176 (1.0231)	N/A

*, **, and *** represent significance on a 90%, 95%, and 99% confidence interval, respectively

[20] Of course, this does not completely isolate the effect of interest from all confounding factors. However, this is the best I can realistically attain given the interconnectedness of the algorithm.

[21] e.g., τ and/or θ on Japanese exports to Canada are usually different from τ and/or θ on Canadian exports to Japan

Overall, 27 out of 42 regression coefficients are significant; of those 27, 22 are positive and only 5 are negative. The average coefficient across all ordered country pairings (where coefficients not significantly different from 0 are counted as 0) is 0.5575, signifying that in general, every 1% increase in shipping price markups results in the satisfaction of PPP conditions being 0.5575 percentage points less likely.

The key limitation of the methodology in this section (5.3) is that, because of the way the algorithm is set up, it is very difficult to isolate the marginal effect of *just* profit-maximizing shipping on PPP failure. There are simply too many confounding factors to measure this exactly (for example, remember that SRPL between some country *A* and another country *B* is influenced by shipping markups on exports from *all* other countries, not just *A* and *B*). Because this section's findings are such a rough estimation, I do not mean for it to be the central finding of this paper. However, I feel that after reading 5.1 and 5.2, many readers will be left asking the question that 5.3 tries to answer. There is no perfect way to answer this question exactly, but this is my best attempt at doing that.

6: Model Limitations and Areas for Future Research

There are several limitations of this algorithmic model worth mentioning. First and foremost, this is not a general equilibrium model, in that it assumes labor and capital markets are in equilibrium exogenously and thus firms' production costs are exogenous; consumer expenditure levels are also determined exogenously. Furthermore, financial markets are not a component of this model in any capacity, not even an exogenous component. I felt that the model is complex enough as it is —and it takes enough time for the program to run as is— that modeling additional external markets would be impractical. However, a paper that gracefully implements a general equilibrium version of this paper's model would be an interesting read.

Second —and this builds off of the fact that this is not a general equilibrium model— firms are static and cannot move from one country to another. In the real world, if shipping rates from some country *A* to some other country *B* are consistently higher than vice versa, after a while producers in *A* who want to sell to *B* will simply move their production operations to *B*, *ceteris paribus*. Along these same lines, there is no firm entry/exit condition in the model, meaning that production of goods is entirely exogeneous. Consumers cannot migrate either, but this is not as large of a concern in this model as firms not being able to migrate.

Third, there is no explicit distinction between intermediate and final goods. This is one particular manifestation of the first limitation, that production costs are exogenous. The KLEMS data counts intermediate input costs as a component of production costs, so intermediate goods are considered in *some* way. However, I assume that all goods are eventually purchased by consumers; thus, consumers are always the ones who determine demand as opposed to other firms.

Fourth, shipping line capacities are assumed infinite; I assume that shipping lines are always able to accommodate additional demand for shipping between a certain pair of countries on a certain day, and these firms choose the profit-maximizing shipping price accordingly. However, if total export demand were ever to exceed that which the shipping line can handle, it would create further inefficiencies in export markets that this model would not capture.

Finally, home biases in trade are not captured. This type of bias, which is well-documented by Obstfeld and Rogoff (2000), Wolf (2000), Yi (2010), and many others, occurs when consumers tend to prefer goods produced in their own country, even if the same good is available for cheaper when imported from elsewhere. Instead, I assume that consumers have no preference for domestic goods, apart from the fact that domestic goods will tend to have depreciated to a lesser extent than foreign goods by the time they reach destination markets; thus, a certain good produced in a foreign country will tend to be of slightly lower quality than a comparable good produced domestically, when the consumer actually buys it.

7: References

7.1: Literature

Akinci, O. (2017). A Note on the Estimation of the Atemporal Elasticity of Substitution Between Tradable and Nontradable Goods. *Federal Reserve Bank of New York*. Retrieved from http://ozgeakinci.com/wp-content/uploads/2017/04/ElasticityofSubstitution_TRvsNTR_.pdf

Alaouze, C. M. (1977). Estimation of the Elasticity of Substitution Between Imported and Domestically Produced Goods Classified at the Input-Output Level of Aggregation. *Impact Project*, *O-13*. Retrieved from http://www.copsmodels.com/ftp/workpapr/o-13.pdf

Anderson, J. E., & Wincoop, E. V. (2004). Trade Costs. *Journal of Economic Literature*, *42*(3), 691–751. doi: 10.1257/0022051042177649

Atkeson, A., & Burstein, A. (2008). Pricing-to-Market, Trade Costs, and International Relative Prices. *American Economic Review*, *98*(5), 1998–2031. doi: 10.21034/sr.404

Bernard, A. B., Jensen, J. B., & Schott, P. K. (2006). Trade costs, firms and productivity. Journal of Monetary Economics, 53(5), 917–937. doi: 10.1016/j.jmoneco.2006.05.001

Bosker, M., & Buringh, E. (2017). City seeds: Geography and the origins of the European city system. *Journal of Urban Economics*, *98*, 139–157. doi: 10.1016/j.jue.2015.09.003

Dixit, A. K., & Stiglitz, J. E. (1977). Monopolistic Competition and Optimum Product Diversity. *American Economic Review*, *67*(3), 297–308. Retrieved from https://www.jstor.org/stable/1831401

Duarte, M., & Stockman, A. (2001). Rational Speculation and Exchange Rates. *NBER Working Paper Series*, *8362*. doi: 10.3386/w8362

Frydman, R., Goldberg, M. D., Johansen, S., & Juselius, K. (2009). A Resolution of the Purchasing Power Parity Puzzle: Imperfect Knowledge and Long Swings. *Aarhus School of Business*. doi: 10.2139/ssrn.1326396

Hummels, D. L. (1999). Toward a Geography of Trade Costs. *SSRN Electronic Journal*. doi: 10.2139/ssrn.160533

Irarrazabal, A., Moxnes, A., & Opromolla, L. D. (2015). The Tip of the Iceberg: A Quantitative Framework for Estimating Trade Costs. *Review of Economics and Statistics*, *97*(4), 777–792. doi: 10.1162/rest_a_00517

Krugman, P. (1995). Increasing returns, Imperfect Competition and the Positive Theory of International Trade. *Handbook of International Economics*, *3*, 1243–1277. doi: 10.1016/s1573-4404(05)80004-8

Krugman, P. (1980). Scale Economies, Product Differentiation, and the Pattern of Trade. *The American Economic Review*, *70*(5), 950–959. doi: 10.7551/mitpress/5933.003.0005

Lu, D. (2018, February). *ECO270: International Economics*. *ECO270: International Economics*. Rochester. Retrieved from https://learn.rochester.edu/bbcswebdav/pid-1333630-dt-content-rid-4955824_1/courses/ECO270.2018SPRING.84469/Lecture5.pdf

Lu, D. (2018, February). *ECO270: International Economics*. *ECO270: International Economics*. Rochester. Retrieved from https://learn.rochester.edu/bbcswebdav/pid-1333630-dt-content-rid-4955824_1/courses/ECO270.2018SPRING.84469/Lecture6.pdf

Melitz, M. J. (2003). The Impact of Trade on Intra-Industry Reallocations and Aggregate Industry Productivity. *Econometrica*, *71*(6), 1695–1725. doi: 10.1111/1468-0262.00467

Miklosovic, T., & Lichner, I. (2011). The Estimation of Elasticity of Substitution of CES Production Function. *SSRN Electronic Journal*. doi: 10.2139/ssrn.1874586

Murray, C. J., & Papell, D. H. (2005). The Purchasing Power Parity Puzzle is Worse Than You Think. *Empirical Economics*, *30*(3), 783–790. doi: 10.1007/s00181-005-0261-9

Obstfeld, M., & Rogoff, K. (2000). The Six Major Puzzles in International Macroeconomics: Is There a Common Cause? *NBER Macroeconomics Annual*, *15*, 339–390. doi: 10.1086/654423

Parsley, D., & Wei, S.-J. (2001). Explaining the Border Effect: The Role of Exchange Rate Variability, Shipping Costs, and Geography. *Journal of International Economics*, *55*(1), 87–105. doi: 10.3386/w7836

Ricci, L. A., & Macdonald, R. (2002). Purchasing Power Parity and New Trade Theory. *IMF Working Papers*, *02*(32). doi: 10.5089/9781451845129.001

Rogoff, K. (1996). The Purchasing Power Parity Puzzle. *Journal of Economic Literature*, *34*(2), 647–668. Retrieved from https://www.jstor.org/stable/2729217?seq=1#metadata_info_tab_contents

Salem, H. H. (2005). The estimation of the elasticity of substitution of a CES production function : Case of Tunisia. *Global Trade Analysis Project*, *1626*. Retrieved from https://www.gtap.agecon.purdue.edu/resources/res_display.asp?RecordID=1626

Samuelson, P. A. (1954). The Transfer Problem and Transport Costs, II: Analysis of Effects of Trade Impediments. *The Economic Journal*, *64*(254), 264–289. doi: 10.2307/2226834

Spelic, J. (2017, August 16). PartnerShip LLC. Retrieved February 17, 2020, from https://www.partnership.com/blog/post/ltl-vs-truckload-freight-what-s-the-difference

Taylor, A., & Taylor, M. (2004). The Purchasing Power Parity Debate. *Journal of Economic Perspectives*, *18*(4), 135–158. doi: 10.3386/w10607

What is an FAK (Freight All Kinds) and How Can It Help Your Rates? (2019, June 10). Redwood Logistics. Retrieved March 10, 2020, from https://www.redwoodlogistics.com/what-fak-how-can-it-help-your-rates/

Wolf, H. C. (2000). Intranational Home Bias in Trade. *Review of Economics and Statistics*, *82*(4), 555–563. doi: 10.1162/003465300559046

Xu, J. (2010). Noise Traders, Exchange Rate Disconnect Puzzle, and the Tobin Tax. *Journal of International Money and Finance*, *29*(2), 336–357. doi: 10.1016/j.jimonfin.2009.06.013

Yi, K.-M. (2010). Can Multistage Production Explain the Home Bias in Trade? *American Economic Review*, *100*(1), 364–393. doi: 10.1257/aer.100.1.364

7.2: Data Sources

Source	**URL**
FRED	https://fred.stlouisfed.org/
Bureau of Economic Analysis (Depreciation)	https://apps.bea.gov/iTable/index_FA.cfm
Freightos Transit Time Calculator	https://www.freightos.com/freight-resources/transit-time-calculator-for-international-freight-free/
Freightos Baltic Index	https://fbx.freightos.com/
KLEMS	http://www.worldklems.net/data.htm
Freight Analysis Framework	https://faf.ornl.gov/fafweb/Extraction0.aspx
Bureau of Transportation Statistics	https://www.bts.gov/browse-statistical-products-and-data
CEIC	https://www.ceicdata.com/en

8: Appendices

Appendix A: Profit-Maximizing Condition for Producers

We can rewrite (4) as an inverse demand function:

(i) $$p_{c,t-d}\left(q_{c,t-d}\right) = A_{c,t}\frac{q_{c,t-d}^{\frac{-1}{\sigma}}}{(1-v_t)^d}$$

with $A_{c,t}$ being an exogenous constant. The profit maximizing condition occurs at the price where $MR = MC$. In our case, I can rewrite the profit maximizing condition $MR = MC$ for each potential export partner c as

(ii) $$p_{c,t-d}\left(q_{c,t-d}\right) + \frac{\partial p_{c,t-dc}\left(q_{c,t-d}\right)}{q_{c,t-d}} = A_{c,t}\frac{q_{c,t-d}^{\frac{-1}{\sigma}}}{(1-v_t)^d} - \frac{1}{\sigma}A_{c,t}\frac{q_{c,t-d}^{-\frac{1}{\sigma}-1}}{(1-v_t)^d}q_{c,t-d} = m_t + \tau_{c,t}\ .$$

Substituting the expression for $p_{c,t}$ from (3) into (4), and then substituting that updated version of (4) into (ii) yields

(iii) $$A_{c,t}^{1-\sigma}\left(1-\frac{1}{\sigma}\right)\left(\frac{p_{c,t}}{(1-v)^d}\right) = m_t + \tau_{c,t}$$

which means

(iv) $$p_{c,t} = \left(\frac{\sigma}{\sigma-1}(1-v)^d\right)\left(\frac{1}{A_{c,t}^{1-\sigma}}\right)\left(m_t + \tau_{c,t}\right).$$

Note that in (5), $K_{c,t}$ is equivalent to $A_{c,t}^{1-\sigma_c}$.

Appendix B: To Export or Not To Export?

Starting with equation (6):

(6) $$Pred\left(p_{c_2,t}q_{c_2,t}\right) - \left(m_{c_1,t-d} + \tau_{c_1c_2,t-d}\right)Pred\left(q_{c_2,t}\right) > 0$$

I first focus on the $Pred\left(p_{c_2,t}q_{c_2,t}\right)$ component. We can substitute the following for $p_{c_2,t}q_{c_2,t}$ in (6), according to equations (2), and then (5):

(v) $$p_{c_2,t}q_{c_2,t} = p_{c_2,t}\left(\frac{p_{c_2,t}^{-\sigma_z}}{A_{c_2,t}}\right) = \frac{p_{c_2,t}^{1-\sigma}}{A_{c_2,t}} = \frac{\left(\frac{\sigma}{\sigma-1}(1-v_t)^d\left(\frac{1}{K_{c_2,t}}\right)\left(m_{c_1,t-d}+\tau_{c_1c_2,t-d}\right)\right)^{1-\sigma}}{A_{c_2,t}}\ .$$

The firm can predict elasticity of substitution σ perfectly because it is time-invariant, so that the firm views it as a given/constant. The firm incurs depreciation costs continuously from day $t-d$ until *t*, so I decide to use v_{t-d} instead of v_t for the sake of mathematical tractability. Additionally, the firm incurs its marginal costs of production and transport costs on day $t-d$, so it wants to set its markups proportional to what they actually incurred; thus, it considers marginal cost of production and transport costs on day $t-d$, each of which the firm already knows and does not have to predict. Building these assumptions into an expanded version of (v) gives me:

(vi) $$Pred\left(p_{c_2,t}q_{c_2,t}\right) = Pred\left(\frac{p_{c_2,t}{}^{1-\sigma}}{A_{c_2,t}}\right) =$$
$$\left(\tfrac{\sigma}{\sigma-1}\right)^{1-\sigma}(1-v_{t-d})^{d-d\sigma}\left(m_{c_1,t-d}+\tau_{c_1c_2,t-d}\right)^{1-\sigma} Pred\left(K_{c_2,t}{}^{\sigma-1-\frac{1}{1-\sigma}}\right).$$

Now for the $\left(m_{c_1,t-d}+\tau_{c_1c_2,t-d}\right)Pred\left(q_{c_2,t}\right)$ component. If I substitute equation (2) for $q_{c_2,t}$, I have:

(vii) $$\left(m_{c_1,t-d}+\tau_{c_1c_2,t-d}\right)Pred\left(q_{c_2,t}\right) = \left(m_{c_1,t-d}+\tau_{c_1c_2,t-d}\right)Pred\left(\frac{p_{c_2,t-d}{}^{-\sigma}}{A_{c_2,t}}\right) =$$
$$\left(m_{c_1,t-d}+\tau_{c_1c_2,t-d}\right)Pred\left(\frac{\left(\frac{\sigma}{\sigma-1}(1-v_t)^d\frac{1}{K_{c_2,t}}\left(m_{c_1,t-d}+\tau_{c_1c_2,t-d}\right)\right)^{-\sigma}}{A_{c_2,t}}\right)$$

which simplifies to:

(viii) $$\left(m_{c_1,t-d}+\tau_{c_1c_2,t-d}\right)^{1-\sigma}\left(\tfrac{\sigma}{\sigma-1}(1-v_t)^d\right)^{-\sigma} Pred\left(K_{c_2,t}{}^{\sigma-\frac{1}{1-\sigma}}\right)$$

when treating known values as constants (which do not have to be predicted). Substituting (vi) and (viii) into (6) yields (7).

Appendix C: Profit-Maximizing Condition for Shipping Lines

Building off of equation (10), the following represents the first-order condition to maximize profits for shipping lines:

(ix) $$\frac{\partial}{\partial\tau_{c_1c_2,t-d}}\sum_{\zeta}\left(\tau_{c_1c_2,t-d}-\theta_{c_1c_2,t-d,\zeta}\right)\frac{\left(\frac{\sigma_\zeta}{\sigma_\zeta-1}\left(\frac{1}{K_{c_2,t-d}}\right)\left(m_{c_1,t-d,\zeta}+\tau_{c_1c_2,t-d}\right)\right)^{-\sigma_\zeta}}{K_{c_2,t-d}} = 0\,.$$

But shipping lines cannot predict which producers will be trying to ship on that day, given that producers make their exporting decisions *in response to*, and not *ahead of*, the prices that shipping lines set. Hence, the shipping line assumes that the producer of every good *z* wants to ship on that day (the set of all ζ's is a subset of all *z*'s). With that in mind, the first order condition that shipping lines really face is:

(x) $$\frac{\partial}{\partial\tau_{c_1c_2,t-d}}\sum_{z}\left(\tau_{c_1c_2,t-d}-\theta_{c_1c_2,t-d,z}\right)\frac{\left(\frac{\sigma_z}{\sigma_z-1}\left(\frac{1}{K_{c_2,t-d}}\right)\left(m_{c_1,t-d,z}+\tau_{c_1c_2,t-d}\right)\right)^{-\sigma_z}}{K_{c_2,t-d}} = 0$$

After taking the derivative by applying the product rule, I have:

(xi)

$$\sum_z \left[\left(\tau_{c_1c_2,t-d} - \theta_{c_1c_2,t-d,z}\right) \frac{-\sigma_z^2}{(\sigma_z-1)K_{c_2,t-d}} \frac{\left(\frac{\sigma_z}{\sigma_z-1}\left(\frac{1}{K_{c_2,t-d}}\right)\left(m_{c_1,t-d,z}+\tau_{c_1c_2,t-d}\right)\right)^{-\sigma_z-1}}{K_{c_2,t-d}} + \frac{\left(\frac{\sigma_z}{\sigma_z-1}\left(\frac{1}{K_{c_2,t-d}}\right)\left(m_{c_1,t-d,z}+\tau_{c_1c_2,t-d}\right)\right)^{-\sigma_z}}{K_{c_2,t-d}} \right] = 0$$

There is no way to solve for $\tau_{c_1c_2,t-d}$, or at least not according to any math I know. Instead, I employ Gradient Ascent Search. This algorithm follows the same concept as Gradient Descent Search. In Gradient Descent Search, one estimates a loss function to evaluate loss for each iteration. If loss decreases from the previous iteration, repeat until there is an iteration in which loss finally increases. When this happens, it means that we have arrived at a local minimum of the loss function. In the case of this project, I use Gradient *Ascent* Search, where instead of trying to minimize a loss function, I try to maximize a profit function based on a chosen price. Of course, this algorithm is sometimes suboptimal when we get “stuck” at a local optimum instead of a global optimum. However, I feel it is the best option available to me given the limitations of my computer’s hardware.