Department for
Business & Trade

# HMG's Partial Equilibrium Trade Models (PE-D)

## 1. Introduction

His Majesty's Government (HMG) is pursuing a number of changes to trade policy.

It is vital that HMG makes full use of all available analytical tools to both inform policy decisions before the event (so-called ex ante analysis) and assess their impacts during and after (ex post analysis). These tools include (but are not limited to) a suite of trade models each with different strengths and areas of focus. This fits with the recommendations of the Trade Modelling Review's Expert Panel[1] and is essential for robust and comprehensive economic analysis in line with the principles of the Government Economic Service.

The Trade Modelling Expert Panel emphasised the importance of complementing the department's core Computable General Equilibrium modelling with an expanded capacity to 'zoom-in' on sensitive or high-impact sectors. As stated by the panel, this necessitates the 'further development of "partial equilibrium" modelling, drawing on a variety of data sources and analytical methods'.

PE models tend to be used in trade policy analysis to assess the potential impact of new policies such as agreeing a Free Trade Agreement (FTA) between two countries, although they can also be used to analyse the impact of past policy changes. They provide a means of testing how impacts might vary depending on the nature of the policy changes and on which changes might be more significant than others. For an FTA they can estimate the impact on the countries directly involved in the FTA and on third countries.

The government has developed a suite of several Partial Equilibrium models. The DBT (2025) working paper[2] details how HMG uses a range partial equilibrium models by presenting results from running simulations of a notional trade agreement between two countries. It uses these to show how these models' function, what they can deliver, what are the key parameters that drive their results and how sensitive the results are to changes in these parameters. One, the PE-Trade or PETRA model, is outlined in another documentation note[3].

Another PE model, described in this paper, is the PE–D model. Partial Equilibrium (PE) trade models are essentially a set of equations based on economic theory that incorporate a range of factors which influence the prices and sales of products, whether imported or produced domestically. PE-D estimates the economic impact of changes in barriers to trade,

---

[1] See https://www.gov.uk/government/publications/trade-modelling-review-expert-panel-report-and-recommendations/trade-modelling-review-expert-panel-report

[2] Working paper illustrating the use of HMG's partial equilibrium trade models (July 2025) https://www.gov.uk/government/publications/partial-equilibrium-trade-models-modelling-paper

[3] HMG's partial equilibrium trade model (PE-TRAde or PETRA) (July 2025) https://www.gov.uk/government/publications/partial-equilibrium-trade-pe-trade-or-petra-model-modelling-paper

such as tariff rates and non-tariff measures. It estimates the changes to a range of variables from their initial values, including domestic production and trade.

Like all partial equilibrium models, PE-D focuses on the direct or ‘first order’ impact of a policy change on a particular sector. It does not incorporate general equilibrium effects that might result from policy changes, for example from a reallocation of resources or changes in capital allocation, relative wages or employment. This makes it easier to see the potential ‘first order’ causes and effects from the policy changes being modelled. Its simpler structure means it is less computationally complex and data intensive than general equilibrium models and also provides greater scope for sectoral disaggregation.

PE models simulate possible impacts resulting from a policy change; they are not a forecast but rather are intended to guide their users to the potential direction of movement and order of magnitude of possible changes as well as how sensitive these might be to variations in the policy changes.

PE-D is ‘static’, which means it simulates the change from an initial equilibrium period, based on historical data, to a new equilibrium once all the impacts of the policy change that are being modelled have worked their way through the sectors in the model. It does not predict the path of how the economy will move to its new equilibrium. Nor does it consider how other factors such as demographics or productivity may change over time.

It is possible that a new equilibrium may be reached in a modelling sense faster in a PE than in a Computable General Equilibrium (CGE) model, such as the GTAP model used by the Department to simulate the potential impact of Free Trade Agreements. We would expect labour and capital markets to adjust more slowly than goods markets, therefore take more time to reach their new equilibrium and, unlike CGE models, PE models do not assume that labour or capital markets adjust.

PE-D is intended to complement the results from other HMG models, especially the CGE model, by being able to simulate potential impacts at a more disaggregated product level. Because it requires less data than PETRA or CGE models, it can be run for more granular commodities. The standard PETRA dataset has around 120 manufacturing and agricultural food sectors but does not include any services, minerals  or plants and animals. PE–D can be run for any commodity for which sufficient data is available. Because it can be used to simulate impacts at the level of HS6 lines, there are potentially thousands of different commodities which could be modelled.

Even though PE-D offers the option of far more granular analysis than PETRA if it is run at an HS6 line level, its sectors can still contain a variety of products with different characteristics and preferences. Therefore, caution should be taken before assuming that the results for each sector necessarily apply to all products. Instead results need to be interpreted bearing in mind the following differences between products contained within the sector. For example, the HS6 line 870322 contains all types of passenger motor cars with petrol engines between 1 and 1.5 litres. This is much more differentiated than the ISIC4 sector 2910, which includes many different vehicles of different engine types and capacities but still can include cars with very different characteristics which will attract different customers. Different countries may have very different patterns of demand for these cars.

Moreover, tariff and non-tariff barriers may vary considerably across the products within a sector, especially for food products. There is thus a risk of aggregation bias where the average barriers for a sector may not be appropriate for all products within the sector, especially if there are significant peak barriers for some products.

Features of PE models which apply to PE-D are that it does not consider how changes in wages, employment, or investment in one sector affect other sectors.

Nor does the model consider how sectors might be affected by changes in the costs of their inputs.

Nor should the results for individual PE products or sectors be aggregated to estimate the “total” impact from a trade shock as such an estimate would exclude all the intra-sectoral impacts, which can be substantial.

All models are highly dependent on the quality of their input data, especially elasticities. This can become more significant when running models at a more granular level where robust data is harder to find. There is also the issue that some data that is needed is not available in a form that can be readily used in models. For example, models need estimates of non-tariff measure (NTM) costs. Modellers generally assume that the impact of NTMs can be represented by an ad-valoren equivalent (AVE), ie a % increase in costs. This is an acknowledged simplification of the role of NTMs, which in extreme cases may prohibit entry instead of providing an addition to costs.

# 2. Main Features of the PE - D Model

## 2.1 Overview

PE - D simulates the direct impact of a trade policy shock such as tariff or non-tariff changes on trade between two (or more) countries or trading blocs by estimating changes in levels of domestic production (only for the country being simulated) and trade (imports and exports).

It is effectively a demand-driven model. The model draws upon the formal structure set out in Krugman (1980), although using a simplified form of this model. This model has several desirable properties, providing a theoretical basis for several facets of trade that occur in practice but are not predicted by some trade theories. These include the existence of intra-industry trade[4] and home bias, such that larger countries typically have a greater share of their home markets. In the model, market shares are determined by consumers maximising their utility and firms their profits, using a constant elasticity of substitution (CES) framework.

As with most models, PE-D starts by estimating (calibrating) the structural parameters of its economic functions (in its case its demand function) using existing trade flows and trade costs (such as tariffs and NTMs). This generates its base, equilibrium state. It then re-optimises demand after trade costs change to generate its simulated outcomes.

Its minimum data requirements are:

- Bilateral trade for the product being analysed between all the countries being modelled;
- Trade costs of importing the product being analysed into the country being shocked – both historical and simulated;

[4] Intra-industry trade is when Country A exports and imports product X from country B and country B does likewise with country A. It frequently occurs, even though it can be seen as inconsistent with certain theories of trade such as comparative advantage and specialisation. It arises because firms in countries A and B produce different varieties of product X and consumers want both varieties.

- Elasticities for the product being shocked, i.e. how demand for the product will change as its price changes;
- Domestic production for the product being analysed in the country being shocked (whilst it can be run without production data, this limits interpretation to changes in trade flows).

### 2.2 Advantages of Granularity

The PE-D model increases the range of policy questions which can be modelled through increased granularity, provided data is available, compared to CGE models and PETRA.

By removing the requirement for production data covering all products and countries in the model, the PE-D model facilitates PE simulations at a more granular level, i.e. HS6 level and below. Such granular data is rare, but can be assembled for some countries, including the UK. This enables modelling of specific products of interest to stakeholders and can generate more accurate estimates as key parameters such as elasticities and tariffs can vary considerably across products aggregated within the ISIC4 sectors used in PETRA. Even if more granular data isn't available for all products, PE–D can be run for those products where it is available and such data is more likely to be available for one country than across all sectors and countries in the PETRA dataset.

This is particularly useful in tariff analysis or for trade defence cases, where the data and policy instruments are often very fine-grained. For example, it may enable analysis of a particular type of steel products. Whilst simulations at the level of the GTAP or ISIC iron and steel sector can be used as ballpark figures, setting the backdrop for policy discussions, they may not provide robust estimates for specific commodities. The PE-D model offers an alternative, wherever there is production data (or a feasible range) in the country being modelled at the HS6 level.

Interpretation of results at this level of detail requires care. In particular, the PE-D model doesn't capture potential substitutability by different commodities and a comprehensive analysis of changing a trade barrier for a particular commodity would ideally consider its potential for replacement by similar products. The exclusion of substitutability is a greater risk when examining more granular products as at a more aggregated level potential substitutes are more likely to be within the same sector. For example a simulation of the ISIC4 "meats" sector would include the potential of substitution from beef to chicken or lamb, whereas a simulation of a HS6 beef line would not.

## 3. Model's Structure

The model equations are drawn from the micro-based theory proposed by Krugman (1980), which extended to trade the industrial organisation framework of Dixit & Stiglitz (1977). In this environment, an industry is composed of a continuum of small firms trading differentiated goods, with demand following a CES structure[5].

This structure means that as utility displays love-for-variety and as firms' products are differentiated, intra-industry trade from diverse sources is rational even where exporters face additional costs. Effectively, consumers will pay these costs to access varieties of a product distinct from those available from domestic producers.

## *3.1.* Consumers

The prime driver of the model is the consumers utility function, which takes the standard CES form used across most academic macroeconomics and trade theory:

$$U = C = \left(\int_0^N c(i)^{\frac{\sigma-1}{\sigma}} di\right)^{\frac{\sigma}{\sigma-1}}.$$

Here, $N$ is the number of firms active in the sector, $c(i)$ is the consumption of variety $i \in [0, N]$, corresponding to a single firm on the continuum, and $\sigma$ is one of two important elasticities, the elasticity of substitution.[6] $C$ is the CES aggregator, which represents total consumption. Note that, following the guiding principle of this theory, firm location is not relevant for consumer preferences.

The other crucial construct needed to derive the central equations of the model is the CES price index. By weighting each variety proportionately to its share in the consumption bundle of an optimising consumer, this index provides a single number which corresponds to the "shadow price" of an increment in utility. It is defined as:

$$P = \left(\int_0^N p(i)^{1-\sigma} di\right)^{\frac{1}{1-\sigma}}$$

where $p(i)$ is the price of variety $i$.

This provides the relationships used to define the model's main equations. See Annex 1 for details of how they are derived. The first key equation is a closed-form expression for consumer demand:

$$c(i) = p(i)^{-\sigma} P^{\sigma} C$$

which defines the curvature over which an increase in the price of $i$, relative to all other prices as tracked by $P$, is mapped onto a level of demand $c(i)$. $\sigma$ matters a great deal in shaping the response to a shift in relative prices; and, if $\frac{p(i)}{P}$ is large, it will take a greater

---

[5] i.e. the elasticity of substitution between different varieties of a good doesn't change as their market share changes.

[6] As per the 'C' in 'CES', it is assumed that the elasticity of substitution (ie the rate at which demand switches between different varieties of a product as their relative price changes) is independent of their market share. This assumption makes the maths much simpler which is why it is commonly used in economics. Formally it is that for any two varieties (i and j), $\sigma = -\frac{d \ln\left(\frac{c_i}{c_j}\right)}{d \ln\left(\frac{p_i}{p_j}\right)}$.

Or in other words, an increase in the ratio of the price $i$ to the price $j$ by $z$ percent will reduce the demand for $i$ relative to $j$ by $\sigma \cdot z$ percent, whatever i and js initial market shares.

absolute increase in the price to make a change in demand than if $\frac{p(i)}{P}$ were small. These features are the keys to understanding any unusual model results.

## 3.2. Firms

An important feature of the PE-D model with its CES demand with differentiated goods is that it effectively incorporates imperfect competition with firms' earning a constant markup over costs, if profit-maximising behaviour holds across firms.

Defining $\kappa(i)$ as all unit costs of any description faced by firm $i$, the model assumes that

$$p(i) = \frac{\sigma}{\sigma - 1} \cdot \kappa(i).$$

Note that the assumption of small firms eliminates the possibility of endogenous markups and strategic pricing; instead the optimising firm ignores the behaviour of its competitors.

Furthermore, the Dixit-Stiglitz specification is fundamentally demand-driven, in that the productive capacity is taken to respond in the long run to meet demand. As such, unit costs remain constant, i.e. $\frac{d\kappa(i)}{dc(i)} = 0$.

For the purposes of simulating alterations in trade costs[7], it is useful to divide $\kappa$ into two constituent parts, reflecting 'pure' trade costs on the one hand and all other factors on the other:

$$p(i_{od}) = (1 + \tau_{od})(1 + \eta_{od})(1 + \epsilon_{od}) \left(\frac{\sigma}{\sigma - 1}\right) \cdot \gamma(i_o)$$

Here, subscripts $o$ and $d$ refer to the exporting country (origin) and the importing country (destination), such that $p(i_{od})$ gives the price set by exporter $i$ producing in country $o$ and selling in the market $d$. $\tau_{od}$ is then the tariff faced by $o$ when exporting to $d$, $\eta_{od}$ captures non-tariff measures, and $\epsilon_{od}$ refers to any other miscellaneous trade costs. All are taken to be expressed in AVE form. $\gamma(i_o)$ is then the (single) unit cost taken to face all domestic producers in $o$ independent of those related to trade.

Rather than explicitly including each category of trade costs for all exporters, the only costs which need to be stated are those liable to be altered in the simulation. Trade costs which remain constant need not be distinguished from those unit costs specific to domestic production in $o$, and can instead be sifted into a generic *bilateral* unit cost, $\gamma(i_{od})$. Calibrating this term to the data absorbs any relevant role unchanging trade costs would play if stated directly.

Altogether, then, cross-border prices are modelled as a combination of the separable trade cost of interest (e.g. tariffs) and the bilateral cost term:

$$p(i_{od}) = (1 + \tau_{od}) \left(\frac{\sigma}{\sigma - 1}\right) \cdot \gamma(i_{od}).$$

[7] "Trade costs" is a flexible term. It can be limited to tariffs or broadened to include non-tariff measures (NTMs) and potentially other costs associated with trade such as transport. We expect that in most simulations it will be used to cover tariffs or tariffs + NTMs

The price index in the multi-country context is then given by:

$$P_d = \left( \sum_o \int_0^{n_{od}} p(i_{od})^{1-\sigma} di_{od} \right)^{\frac{1}{1-\sigma}} = \left( \sum_o \int_0^{n_{od}} \left( (1+\tau_{od}) \left( \frac{\sigma}{\sigma - 1} \right) \cdot \gamma(i_{od}) \right)^{1-\sigma} di_{od} \right)^{\frac{1}{1-\sigma}}$$

where $n_{od}$ is simply the measure of firms in $o$ participating in the market of $d$. On this basis, counterfactual trade flows can be derived from the consumer demand equation above.

## 3.3. Aggregate Demand

The final ingredient driving a simulation is the overall demand response to an alteration in the price index of the product/sector. Essentially, a further parameter is needed to imitate the general equilibrium re-allocation of aggregate demand following price changes. In PE–D this is captured by a price elasticity of demand $\mu$, such that within a single market

$$\int_0^N c(i)di = bP^{-\mu}.$$

where $b$ is a constant deployed for calibration purposes, meaning aggregate consumption follows price changes according to:

$$\frac{d \ln( \int_0^N c(i)di)}{d \ln P} = -\mu.$$

In other words, a $z$ percent change in the price index corresponds to a $z \cdot \mu$ percent change in the quantity consumed. As a useful reference point, note that $\mu = 1$ delivers a benchmark Cobb-Douglas economy, where quantities adjust to price changes to keep value expenditure constant.

In general, the crucial elasticities $\mu$ and $\sigma$ play different roles, effectively corresponding to an income and a substitution effect respectively. A decline in the price of some varieties of goods in the sector will increase demand for varieties which see no change in price, due to the reduction in $P$.[8] However, in almost all circumstances, this ‘aggregate demand externality’ will be outweighed by downward pressure from the substitution effect, as firms not experiencing a reduction in trade costs will find their loss of market share dominates their gain from the increased size of the market. This is not guaranteed, however, and scenarios can arise where a reduction in the tariff imposed on one nation *increases* demand for its competitors’ products. This is a feature of all CES demand systems and is not unique to the PE-D model.

## 3.4. Small-Shares Adjustment

An additional feature of PE-D which can have a significant impact on its results is the small-shares adjustment (SSA) mechanism, which is described formally in Annex III. This is an option in PE–D, provided that certain conditions, such as imports from the partner country being less than 1% of historic imports in the shocked country, are met.

[8] Perhaps a sudden decline in the price of road bikes permits you to realise your dream of owning both a road bike *and* a mountain bike.

The SSA attempts to fix a known issue with trade models – their difficulty in simulating effects when historical trade is minimal or zero due to prohibitive barriers or missing data (as in such cases trade flows are often recorded as zero). In such circumstances a large shock whilst it could generate a large percentage change in imports, would only have a limited impact on the value of imports, which may be unrealistic. A SSA is especially pertinent in the modelling of more granular sectors, as the possibility of a sector having minimal imports becomes more likely.

PE-D's SSA works by replacing the part of the constant elasticity of substitution (CES) demand curve where imports are close to zero. It changes this to a proxy demand curve based on the primary country's total market as a share of world trade and the partner country's global exports as a share of world trade to set the parameters for the SSA.

Some key points to keep in mind when interpreting the model outputs for 'small-shares' sectors are:

- The country's global export share, will determine its scope for obtaining a positive SSA adjustment; in other words if a country doesn't export much of a commodity to the world, then applying the SSA will have little impact on its exports of the commodity to the shocked country;
- The elasticity of substitution $\sigma$ will dominate the degree of impact the SSA has;[9] and
- Results for sectors where the SSA mechanism are used can diverge significantly from those from running PETRA, or other models without a SSA mechanism.

## *3.5.* Key Inputs: Elasticities

As mentioned above, the elasticities σ and μ are core drivers of the model's results. As a result, selecting the elasticities to use and performing appropriate sensitivity testing around these estimates are pivotal to using the model appropriately.

Unfortunately, they are difficult to estimate. There are a wide range of estimates of σ in the literature, which can be estimated by different, equally justified methods. Whilst there is more agreement about the method for estimating μ, the lack of accessible data on domestic price and quantity variation means it is difficult to estimate in practice.

The following sets out the elasticity estimates currently used in the PE - D model.

Sigma, σ

As discussed earlier, this parameter is responsible for determining the sensitivity of market shares to relative prices, or the extent to which a relative cheapening of one exporter's goods will permit that exporter to displace its competitors in the importing market.

The preferred option is to use econometrically estimated values derived by Fontagné et al. (2022), which provide values down to the HS6 level, derived from a gravity regression. Two rounds of adjustment have been made to these raw elasticities. First, where the estimates

[9] As $\sigma$ gets very low, the small share adjustment mechanism's effect will grow weaker; eventually this will cause the model to crash.

are not statistically significant, they are replaced with the UKTPO supplied PETRA priors, either 3 or 6. Second, the estimates are capped above and below at 2 and 9. This is not to suggest that the elasticities of HS6-level commodities might not exceed these levels. But it has been decided to exclude such extreme values to ensure consistency with the wider suite of HMG trade models.

For sensitivity testing, for the statistically significant estimates, upper and lower bounds are provided to give a 95% confidence interval around the point estimates. For those estimates which have been bound at 2 or 9, the standard errors are re-scaled to remain in proportion with the point estimate, and the confidence interval is constructed on this basis. For insignificant estimates, the UKTPO supplied PETRA priors are perturbed by 50% each way.

All artificially imposed integer values are either increased or decreased by 0.1. This is unlikely to significantly affect results and is generally helpful to avoid using integers given their capacity to produce 'special case' versions of the core functions, which diverge from normal behaviour, particularly when deploying the small-shares adjustment mechanism.

Mu, μ

This parameter measures the degree to which overall demand for a product is sensitive to the average price of the product – it can be thought of as the income elasticity of demand. The estimates suggested for the PE-D model follow PETRA in using import elasticities of demand, derived from the translog GDP method set out by Kee et al. (2008). The actual numbers used come from Ghodsi et al. (2016), who provide HS6-level estimates across a wide range of countries. Where significant, values are provided for both the UK estimate and a 'global' estimate, which is the arithmetic mean of all significant estimates provided. All μ point estimates are bounded between 0.5 and 1.5. Upper and lower bounds of the 95% confidence interval are again provided, with standard errors re-scaled for artificially bounded terms as with the elasticity of substitution estimates.

Whilst treating import elasticities of demand as proxies for overall price elasticities of demand is not ideal, it is preferable to fixing μ as a constant.

For both elasticities, it is recommended to run sensitivity tests, using the maximum and minimum values of the elasticities as well as the central estimate to construct a range of results.

## *3.6.* Other Data and Sources

In addition to the elasticities, PE–D requires data for the size of trade barriers (baseline and simulation); and directly observable data for trade, production and domestic consumption of domestic production (DCDP).

Trade data is readily available and is typically drawn from UN Comtrade (down to HS6) or ITC Trade Map (down to the tariff line). There can be issues around missing data or when different countries have different values for a trade flow (when country A's exports to country B do not match country B's imports from country A, once differences in valuation such as cif and fob are taken into account). Typically, in such cases the average of both countries data is used, although if only one country records a flow it is used instead.

Granular production data is more difficult. Part of the purpose of the PE-D model is to facilitate the use of alternative data if it is available for the importing country's domestic production. For example, for the UK, some manufacturing production data can be obtained from Prodcom and agricultural production data from the annual Defra 'Agriculture in the United Kingdom' reports. Agricultural production (and import data) data in quantities is also available for countries for specific products from FAO Stat. These can be used to create an import penetration multiplier, for instance, 'country X's production of product Y is n times its imports of product Y'. This multiplier can then be applied to trade data in value terms to estimate a value of domestic production consistent with trade figures. DCDP can be calculated by subtracting exports from total production.

More generally, a wider range of non-official sources can be drawn upon such as trade association or industry numbers. A lack of compatible data for partner countries is not a constraint when running the PE-D model.

Baseline trade barrier costs are historical tariffs (plus possibly NTMs rates – especially if these are also changed in the simulation). Tariff rates are typically taken from WITS or MACMAPS. Specific and other non-AVE tariffs need to be converted into ad valorum equivalents (AVEs). NTMs are taken from those in the PETRA dataset or gravity/econometrically derived rates can be used if these are available.

## *3.7.* Values and Quantities

One issue with the PE–D model, is that like most other trade models, it is primarily designed to work with quantities – it simulates the percentage change in quantities for a given shock. But much of the data that it uses is available in values. To tackle this issue, the model has an option to work in terms of "values" or in terms of "quantities".

When using input data in post-tax values (as most trade data will be – as it incorporates tariffs and NTMs), a decision must be made on how to present and interpret the results. If the "quantities" option is selected, the model will take an input value of (for example) £1 million to represent 'the number of units which can be purchased in the baseline for £1 million'. The model does not know what these units are or how many of them are being traded, but the idea is that they are some non-monetary, widget-like object which can be defined.

Following the price shock, widget-demand will change, and using the quantity specification, the output is given in terms of the number of widgets that could have purchased for £1 million[10] *in the baseline*. This is the simplest way to think of an increase in physical production – it is the increase in the value of production at constant prices. It is challenging to interpret, however, because the purpose of the simulation is to *alter* prices by adjusting trade costs. These constant price outputs would therefore never appear in actual trade data and therefore have to be appropriately caveated when used.

[10] Millions of pounds are used here only as an example; the model can be run with values expressed in any currency desired.

The alternative is to run the model on “values”. In this case, both model inputs and outputs can be interpreted as the numbers found in trade data, which gives them a surer ontological footing. To provide these results, the model simply deflates by the price adjustment; it delivers £1 million of widgets in old money, and then calculates what those widgets will be sold for in the new equilibrium.

However, this can generate odd looking raw trade numbers. To take an extreme example, if there was a baseline tariff of 100%, then £1 million of exports in post-tax trade data corresponds to £0.5 million of pre-tax export value, and £0.5 million collected as duties by the importing country. Only the former number really matters to the exporter. Suppose, for example, that full liberalisation of a particularly price-insensitive product only raised the simulated pre-tax export values to £0.75 million. Then, the overall post-tax value would *fall* after the liberalisation (from £1m to £0.75m), as duties go to zero. But the point relevant to producers, ie the value of domestic production, would have risen by 50% (from £0.5m to £0.75m). To address this potential confusion, when the values option is selected, both the baseline and simulation values are split between post-tax trade, pre-tax trade and duties to enable a clearer picture of the impact on producers.

### *3.8.* Non-Production PE-D

In some cases, DCDP data may not be available in the form required. Simulations can still be run in such circumstances, provided certain limitations are accepted. For example, it is usually possible to provide an upper bound of DCDP, even if that is simply the entirety of a more aggregated sector taken from the PETRA dataset. Especially if the main interest is exploring the imapct on trade, it is possible to use such approximations, relying on the fact that DCDP typically only has a small impact on trade flows. Given a lower bound of zero and some upper bound for the sector, it is possible to run two simulations and provide the results as a range.  But such a range should not be confused with the confidence interval constructed in sensitivity testing.

## *4.* Interpreting Results

As PE–D, like PETRA, only simulates the direct impact of a change in trade costs, this makes it simpler to interpret its results than in a general equilibrium model where many more factors influence the reported impacts.

As mentioned above, the elasticities $\mu$ and $\sigma$ play a critical role in driving PE-Ds results. They have  different roles and can be interpreted as corresponding to income ($\mu$) and substitution ($\sigma$) effects.

- A decline in the price of some varieties of a product will reduce P, so increasing demand for all varieties. This is the income effect.
- But in almost all circumstances, this will be outweighed by a substitution effect, as firms not experiencing a reduction in trade costs will find their relative price increases.

The other relevant factors that drive the results are:

- Size of historic consumption and trade. The greater the historical relationship, the greater will be the expected imapct of a shock.
- Size of shock. The larger the change in relative prices as a result of the shock, the larger will be the impact of the shock.

The combination of these three factors – the elasticities, size of historic consumption and shock should explain the results.

The PE-D model can be used to tackle idiosyncratic problems, for example exploring the impact of a shock on a handful of products rather than an economy-wide simulation. In tackling variegated questions on a case-by-case basis, the model should form only one part of the analysis. Nothing which comes out of the model is valid on its own account but should be interpreted in context of knowledge of the sector and how it is expected to respond to shocks.

# 5. Using the PE-D model – Issues

To protect the quality of results, it is necessary to use robust data and to take care when specifying scenarios and interpreting results. There are several issues which should be considered when running the model.

## 5.1. Sector Definitions

PE modelling requires sectors of the economy be treated as independent. Under any circumstances, this generates issues, which can become increasingly acute as simulations become more granular. PE necessitates the choice between treating, say, 'Meat' as a single sector, or 'Beef', or only 'Frozen Beef'. Once sectors are chosen, the competitiveness of different producers in a given sector are calibrated to match the trade and production data under the elasticities provided. But no account is taken of other sectors. There is no possibility of substitution if a Beef sector is chosen between beef or chicken or if a Frozen Beef sector is being modelled between frozen and chilled beef. Sector partitions thus introduce important non-linearities. Even if a Beef sector were composed only of Fresh Beef and Frozen Beef, results from a simulation on the former will not resemble the sum of separate simulations of the latter two.

## 5.2. Use to model Bilateral/Regional shocks

The current version of the PE-D model should not be used to simulate the impact of a trade shock with the rest of the world or multilateral/plurilateral shocks unless all the countries that are affected can be included in the dataset. This is because the model does not include a facility for varying the trade cost/tariff rate for the rest of the world; it can only vary trade costs for countries in the dataset.

## 5.3. Flexible Supply

The PE-D model does not have a facility for restricting supply. It assumes that supply will adjust to match the change in demand. Whilst generally plausible in the long run, there are times when this might not occur. Whilst the standard PETRA dataset assumes supply is elastic and hence does not act as a significant constraint on the simulation equilibrium, it offers the possibility of varying this assumption so that supply constraints can be modelled.

### 5.4. Competitive Structure

The PE-D model assumes Imperfect Competition and product differentiation. It may therefore be inappropriate for simulating homogeneous products in competitive markets, where the Armington (or Perfect Competition versions) of PETRA may be more suitable. This may be particularly an issue for agri-food products.

### 5.5. Supply Chains

The PE-D model doesn't include an option to simulate the impact of a change in the price of imported intermediate inputs as their trade costs change (supply chain effects) that is available in the Imperfect Competition version of PETRA.

## 6. Comparison with PETRA

The PE-D model has a similar structure to the PETRA Imperfect Competition (IC) model, and as such it generally delivers similar results. There are differences between the PETRA IC and Armington versions and the PE-D model shares these differences with respect to the latter as a result of its closeness to the former.

The similarity with PETRA IC has been tested by running a simple scenario involving a reduction in tariffs through both models and comparing their results. For this exercise, the same input data - trade, production, trade costs and elasticities and the same shocks were used in both models.

There was little difference between the results from the PE-D and PETRA models. The PE-D model tracks the PETRA IC model closely, with only frictional differences and no systematic tendency to deviate in a particular direction.

There were however some significant differences between the two models results when the SSA was applied in PE–D.

## 7. Comparison with CGE

There are several reasons why PE-D results may differ from CGE results. They do not include any general equilibrium effects, and as such represent a partial story as they do not reflect possible changes in wages, employment, capital allocation, etc that will affect consumers and producers. Nor do the PE-D results incorporate any possible supply redirection effects. They also will typically be at a more granular level and so may avoid possible aggregation bias. For these reasons PE-D results would not be expected to sum to the CGE results for a sector.

# References


Dixit, A., and Stiglitz, J. 1977. Monopolistic Competition and Optimum Product Diversity. American Economic Review, 67(3): 297-308

Fontagné, L., Guimbard, H., and Orefice, G. 2022. Tariff-Based Product-Level Trade Elasticities. Journal of International Economics, 137

Ghodsi, M., Grubler, J., and Stehrer, R. 2016. Import Demand Elasticities Revisited. wiiw Working Paper

Kee, H. L., Nicita, A., and Olarreaga, M. 2008. Import Demand Elasticities and Trade Distortions. Review of Economics and Statistics, 90(4): 666-682

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

# Annex

## I. Derivations

Section II above provides the core equations needed to *understand and interpret* the PE-D PE model.

This section derives those core results from the micro-foundations of CES utility theory, with atomic firms making differentiated products. Definitions of notations and several blocks of text are repeated from Section II, to avoid the need to switch between Section II and this Annex.

Note that, in general equilibrium, firms' markups are kept fixed by endogenous entry and exit given fixed costs. In the partial equilibrium setting described here, the number of firms (equal to the number of product varieties) is held fixed. As such, the standard presentation specifying fixed costs is omitted.

### (i) *Basics of Dixit-Stiglitz in one country*

Consumers display CES preferences across a continuum of products:

$$U = C = \left(\int_0^N c(i)^{\frac{\sigma-1}{\sigma}} di\right)^{\frac{\sigma}{\sigma-1}}$$

Where $c(i)$ is consumption of a particular variety $i$ and $N$ is the measure of varieties, equal to the measure of firms (one product per firm). $\sigma$ is the elasticity of substitution between goods in the industry. The optimal allocation of income across goods follows from setting up an expenditure-minimization problem:

$$\min E(U) \quad s.t. \quad \left(\int_0^N c(i)^{\frac{\sigma-1}{\sigma}} di\right)^{\frac{\sigma}{\sigma-1}} \geq U.$$

Forming the Lagrangian

$$\mathcal{L} = \int_0^N p(i)c(i)\, di - \lambda\left(\left(\int_0^N c(i)^{\frac{\sigma-1}{\sigma}} di\right)^{\frac{\sigma}{\sigma-1}} - U\right)$$

and setting the derivative with respect to a particular good to zero gives the following first-order condition:

$$0 = p(i) - \lambda\left(\frac{\sigma}{\sigma-1}\left(\int_0^N c(i)^{\frac{\sigma-1}{\sigma}} di\right)^{\frac{1}{\sigma-1}} \cdot \frac{\sigma-1}{\sigma} c(i)^{-\frac{1}{\sigma}}\right)$$

or

$$p(i) = \lambda \cdot C^{\frac{1}{\sigma}} \cdot c(i)^{-\frac{1}{\sigma}}.$$

Raising both sides to the power $1-\sigma$:

$$p(i)^{1-\sigma} = \lambda^{1-\sigma} \cdot C^{\frac{1-\sigma}{\sigma}} \cdot c(i)^{\frac{\sigma-1}{\sigma}}$$

and integrating each side across the continuum:

$$\int_0^N p(i)^{1-\sigma}\, di = \lambda^{1-\sigma} \cdot C^{\frac{1-\sigma}{\sigma}} \cdot \left(\int_0^N c(i)^{\frac{\sigma-1}{\sigma}}\, di\right) = \lambda^{1-\sigma} \cdot C^{\frac{1-\sigma}{\sigma}} \cdot C^{\frac{\sigma-1}{\sigma}} = \lambda^{1-\sigma}.$$

This equation pins down the price index, characterized as the shadow price of a relaxation of the budget constraint. Define this as $\left(\int_0^N p(i)^{1-\sigma}\, di\right)^{\frac{1}{1-\sigma}} \equiv P$.

Plugging it back into the equation above returns the cost-minimizing demand for any given variety:

$$p(i) = PC^{\frac{1}{\sigma}}c(i)^{-\frac{1}{\sigma}} \rightarrow \frac{c(i)}{C} = \left(\frac{p(i)}{P}\right)^{-\sigma} \rightarrow c(i) = p(i)^{-\sigma}P^{\sigma}C$$

which is the core demand equation in Section IV of the main text.

(ii) *Trade Context*

In the PE-D application of this framework, nothing changes except that firms in different locations will arrive at different representative prices, given both explicit and implicit trade costs, and country-specific determinants of productivity.

Total demand in a country $d$ (for defensive) for products from country $o$ (for offensive) follows from individual firm demands just derived:

$$c_{od} = \int_0^{n_{od}} \left(\frac{p(i_{od})}{P_d}\right)^{-\sigma} \cdot C_d\, di_{od} = C_d P_d^{\sigma} n_{od} \bar{p}_{od}^{-\sigma}.$$

Where $n_{od}$ is the measure of varieties moving from $o$ to $d$, $C_d$ is the CES combination of products from all locations consumed in $d$, and the last term follows from the constant value $p(i_{od}) \equiv \bar{p}_{od}$.

Given this demand system, the impact of tariff alterations on trade flows is determined entirely by the price system in the defensive market, which is characterised by the profit maximisation of firms. As firms produce unique products, each is able to command a degree of market power indexed by the elasticity of substitution between goods in the sector, $\sigma$.

A firm $i$'s profits are given by

$$\pi(i) = p(i)c(i) - \gamma(i)c(i)$$

Where $\gamma$ is unit cost and $c$ is quantity demanded and produced. Zero marginal profit[11] entails the first order condition:

$$\frac{\partial \pi(i)}{\partial p(i)} = c(i) + p(i) \cdot \frac{\partial c(i)}{\partial p(i)} - \gamma(i) \cdot \frac{\partial c(i)}{\partial p(i)} = 0^{12}$$

Rearranging for $p$:

[11] As mentioned above, this condition is underwritten in general equilibrium by endogenous entry and exit but is simply enforced here as part of the partial equilibrium scenery.

[12] Crucially, the PE-D framework is a demand-driven model of trade, so that in the long run $\frac{\partial \gamma(i)}{\partial c(i)} \cdot \frac{\partial c(i)}{\partial p(i)} = 0$.

$$p(i) = \gamma(i) - \frac{c(i)}{\left(\frac{\partial c(i)}{\partial p(i)}\right)}.$$

As specified in the previous section, consumer optimization entails the following demand for a particular product:

$$c(i) = C\left(\frac{p(i)}{P}\right)^{-\sigma} = Cp(i)^{-\sigma}P^{\sigma}$$

This gives the required derivative:

$$\frac{\partial c(i)}{\partial p(i)} = -\sigma C P^{\sigma} p(i)^{-\sigma-1}$$

such that prices follow

$$p(i) = \gamma(i) - \frac{Cp(i)^{-\sigma}P^{\sigma}}{-\sigma C P^{\sigma} p(i)^{-\sigma-1}} = \gamma(i) - \frac{p(i)}{\sigma}$$

or

$$p(i) = \frac{\sigma}{\sigma - 1} \cdot \gamma(i)$$

meaning prices reduce to a markup over costs determined by substitutability between varieties, independent of the scale of production and the price decisions of other firms.

Expanding to the multiple country case, pricing behaviour must be compounded with trade costs distinguished by origin (offensive?):

$$p(i_{od}) = (1 + \tau_{od})(1 + \eta_{od})(1 + \epsilon_{od})\left(\frac{\sigma}{\sigma - 1}\right) \cdot \gamma(i_o)$$

Where $\tau_{od}$ is the tariff imposed on goods sold in $d$ originating in $o$, $\eta_{od}$ denotes NTMs, and $\epsilon_{od}$ refers to other miscellaneous trade costs. Costs not altered in the simulation do not have to be distinguished from productivity differences in computing equilibrium for any given domestic market. As such, cross-border prices are modelled as a combination of the separable trade cost of interest (e.g. tariffs) and a generic bilateral competitiveness term:

$$p(i_{od}) = (1 + \tau_{od})\left(\frac{\sigma}{\sigma - 1}\right) \cdot \gamma(i_{od})$$

The price index in the multi-country context is then given by:

$$P_d = \left(\sum_o \int_0^{n_{od}} p(i_{od})^{1-\sigma} di_{od}\right)^{\frac{1}{1-\sigma}} = \left(\sum_o \int_0^{n_{od}} \left((1 + \tau_{od})\left(\frac{\sigma}{\sigma - 1}\right) \cdot \gamma(i_{od})\right)^{1-\sigma} di_{od}\right)^{\frac{1}{1-\sigma}}$$

on the basis of which counterfactual trade flows can be derived.

## II. Calibration and Solution

The model begins by calibrating the bilateral competitiveness terms based on $\sigma$, the defensive country's import data, and the defensive country's domestic consumption of domestic production. $\gamma_{wd}$, the bilateral unit cost of RoW exporters selling to $d$, is normalized to one.[13] This permits inference to the other bilateral productivity terms by way of the observed trade flows.[14] Noting that

$$\frac{c_{od}}{c_{wd}} = \frac{n_{od}p_{od}^{-\sigma}P_d^{\sigma}C}{n_{wd}p_{wd}^{-\sigma}P_d^{\sigma}C} = \frac{n_{od}}{n_{wd}} \cdot \left(\frac{p_{od}}{p_{wd}}\right)^{-\sigma}$$

and

$$p_{wd} = \left(\frac{\sigma}{\sigma-1}\right)\gamma_{wd} = \left(\frac{\sigma}{\sigma-1}\right)\cdot 1^{15}$$

$$p_{od} = (1+\tau_{od})\left(\frac{\sigma}{\sigma-1}\right)\gamma_{od}$$

this can be rearranged to give $\gamma_{od}$ as a product of observed trade flows, tariffs, fixed $n$ terms and normalized bilateral RoW unit cost:

$$\gamma_{od} = \frac{1}{(1+\tau_{od})} \cdot \left(\frac{n_{wd}c_{od}}{n_{od}c_{wd}}\right)^{-\frac{1}{\sigma}}.$$

Prices then follow from these bilateral unit cost terms, and the price index comes analytically. To complete the baseline calibration, a value must be attached to a multiplier matching demand as a function of the price index to observed total demand:

$$\sum_o c_{od} = b_d P_d^{-\mu} \to b_d = P_d^{\mu} \sum_o c_{od}.$$

This characterises the full set of structural parameters, which remain constant underpinning simulations.

Simulation then introduces a new value for the detachable trade cost component, given these fixed terms. This permits the computation of a new set of prices and a new price index. The price index entails a new value of total demand given the calibrated value of $b_d$ above. The prices themselves, coupled with the fixed $n$ terms, are a sufficient statistic for the relative demand for products from each country. To scale these relative terms to the value of total demand, a new coefficient must be introduced:

$$\alpha = \frac{b_d P_d^{-\mu}}{(\sum_o n_{od}p_{od}^{-\sigma})}$$

Where $P_d$ refers to the newly calculated price index. This coefficient scales the relative demand terms to give new equilibrium trade flows for the market in question:

$$c_{od} = \alpha \cdot n_{od}p_{od}^{-\sigma}.$$

The model then delivers an output data frame containing the new demand from each source in each market assessed.

## III. Small-Shares Adjustment

[13] This is used as the only $\gamma$ term which is always finite.

[14] For brevity, DCDP is understood in this section as bilateral trade flow country $d$ exports to itself.

[15] There is no explicit tariff on imports from the rest of world, as these are absorbed into the definition of $\gamma_{wd}$.

As is standard in international trade models, the PE-D model faces a dilemma in simulating extensive margin adjustment, that is, in simulating trade flows from a starting point of zero or negligible trade. The basic (desirable) property of CES demand is that it works with percentages – a 1% increase in price entails *ceteris paribus* a 1% decrease in demand. An implication of this is that it would require an infinitely large price increase to drive imports of a good to zero, and as such, products with zero imports in the baseline cannot naturally move to positive trade whatever the size of the finite barrier reduction.

More formally, recall that CES demand works in terms of percentages, with a small value of $\frac{p(i)}{P}$ requiring a greater reduction in $p(i)$ to deliver an absolute gain in $c(i)$ than is needed to realise the same gain starting from a large $\frac{p(i)}{P}$. Naturally, there are many zeros in the matrix of bilateral trade flows. To understand the small-shares problem, it is sufficient to ask how large $\frac{p(i)}{P}$ would have to grow before $c(i)$ would converge to deliver these zeros.

From central pricing equation above, it is clear that with $\sigma$ and $\tau_{od}$ fixed from the data, it holds that

$$\lim_{p(i_{od})\to\infty} \gamma(i_{od}) = \infty.$$

Even with a sizeable tariff reduction, there is no coming back from this for the aspiring exporter. This is an appropriate way of treating countries which simply do not export a given product; but there are many cases where a competitive global exporter is entirely unable to penetrate a particular market under extensive trade costs but would be expected to if the cost were removed. This cannot be simulated by any isoelastic demand function, for which such outcomes are unintelligible.

As the number of zeros grows with the granularity of modelling, this limitation occurs more frequently. Hence the need for a fix that can be imposed when market shares are zero or close to zero. The PE-D model deploys an adjustment mechanism, where a country's global export share determines its capacity to 'break into' a market by overriding the behaviour of the CES demand curve at the far right tail.

The central issue to devising a 'good' small shares fix is that whilst we are happy with the behaviour of demand (ie how demand reacts to a change in price) over the central portion of its function we are unhappy with its behaviour in the tails. The objective is therefore to extend the interior (central portion) of the function out to catch extreme cases. This is straightforward, given that the interior of the function is governed by lower-order properties, and behaviour at the extremities is driven by higher-order. The method simply requires approximating the function by perturbing an interior point outward with a lower-order approximation, thereby capturing the desirable part of the function, and dropping the part which is objectionable.

The following method is used:

- The model is calibrated based on existing trade. Cases of negligible imports which might benefit from having the SSA applied are identified by applying the following criteria:

- Current rule of thumb is sector/countries with <1% of the domestic market,
    - where the exporting country's exports (are >1% of the world total),
    - when the importing country has significant imports in the sector.
- If extensive margin adjustment is appropriate given these checks, the calibrated values of $C$ and $P$, the CES aggregation of existing consumption and the price index, are taken.
- For the purposes of small-shares calibration, these are held fixed, so that the export equation (see Section IV) becomes a one-dimensional demand curve: $c_{od} = (n_{od}C_d P_d^{\sigma}) \cdot p_{od}^{-\sigma}$, giving $c_{od}$ in terms of $p_{od}$ with constant bracketed terms.
- $p_{od}$ is then adjusted to drive $c_{od}$ to equal a share of the defensive market's imports equivalent to the exporter's share of global exports in the sector, to capture the 'average exports' case.
- A third-order Taylor expansion is taken around the value of $p_{od}$ which attains this; the resulting polynomial is solved for a value of $p_{od}$ setting it to zero (unlike the CES function, this will be a finite and reasonable number).
- This $p_{od}$ is decomposed into the explicit trade cost, markup, and competitiveness parameter $\gamma_{od}$. The $\gamma_{od}$ term is taken as a new structural parameter, and the originally calibrated equilibrium is shocked to incorporate it, resulting in positive imports from the adjusted partner. This provides a new baseline, on which the actual trade shock is performed.

The method extends behaviour around the average export share of a country to the point of zero exports, and then converts this back into a structural parameter intelligible in the CES framework. It should be noted that the new baseline produced in this way is only meaningful as part of the general shock – in the new baseline equilibrium, trade flows will be altered from their real-world values. The justification for this 'double' impact stems from the fact that there is both the smooth price-adjustment component of price response, and the discontinuity implied by moving from infeasible to feasible trade. The jump from third-degree Taylor polynomial to CES curve thus targets an actual component of a trade shock capable of inducing extensive margin adjustment and is not a distortion in principle.

One important caveat to this approach is that the $\gamma$ term produced is the *lowest* unit cost consistent with taking the Taylor polynomial to zero – it begins from the assumption that imports are 'only just' zero. There is no clear alternative to this, as an observed lack of trade cannot provide any information on just how far removed the producer is from attaining positive exports.

*With $C$ and $P$ fixed, the solid green curve plots the one-dimensional CES demand function, the dash-dot purple curve plots the third-degree Taylor polynomial, the horizontal line gives zero exports, the left blue vertical line gives the price corresponding to a country achieving its average export share in the importing market, and the right red vertical line gives the price corresponding to 1% of the importing market, as a rough idea of the edge of the interior of the function and the beginning of the problematic tail region. The objective of the approximation used is that error should be small within the interior, and large in the small-share section, which is generally attained reasonably well; note that a first-order derivative would clearly induce major error throughout the interior.*

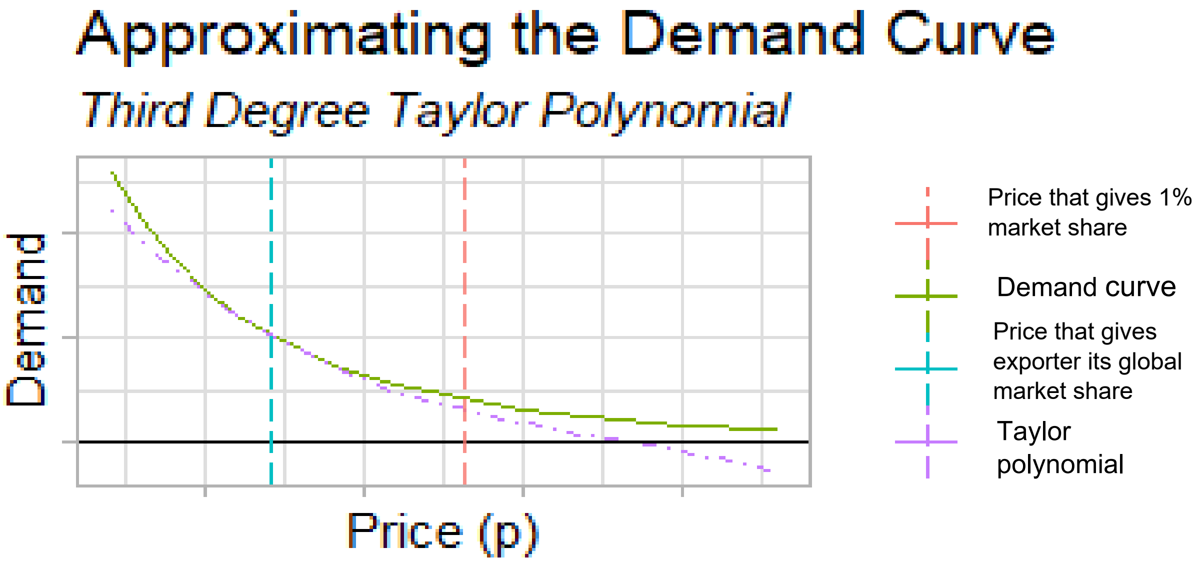


This diagram illustrates the process outlined in the box on the previous page.

The green curve is the model's CES demand curve, with C and P fixed, whose slope becomes very flat as the price rises and demand approaches zero, so much so that it never actually reaches zero (crosses the x axis).

The dot-dash purple curve represents the Taylor polynomial around the demand curve. In this diagram it is shown as having the desired properties of: a steeper slope in the zone in which the Small Shares Adjustment (SSA) may be applied (between the vertical red line and where it crosses the Price axis) and a similar slope to the demand curve around the price which would leave the exporter with its global market share (where it crosses the vertical blue line).

The red vertical line represents the Price which generates a 1% share of the market being shocked – one of the triggers of the SSA.

The blue vertical line represents the Price which would provide the exporter with its global market share in the market being shocked.