Inequality

The IFS Deaton Review

Trade and price-index inequality

David Atkin

An IFS initiative funded by the Nuffield Foundation

IFS Institute for Fiscal Studies

Nuffield Foundation

Trade and price-index inequality

David Atkin (MIT; NBER; CEPR; BREAD)[1]

Introduction

David Dorn and Peter Levell's chapter (Dorn and Levell, 2021) provides a broad and clear overview of the relationship between trade and inequality for developed countries such as the United Kingdom. Trade may affect inequality by making either nominal incomes or price indices less equal. While the impacts of import competition on employment and wages have garnered substantial interest in recent years,[2] two issues have received less attention in the literature: the gains from trade accruing to the very top of the income distribution and the effect of trade on the price index.

The effect of trade on the incomes of the very rich not only affects inequality directly but may also affect the progressiveness of government policy given the out-sized political influence of this group. Our knowledge gap comes in part from the difficulty identifying clean export shocks for developed countries (with the impact of Eastern European integration to the European Union on Germany being perhaps an exception). Even then, the big winners are likely to be investors, owners, managers and highly skilled employees who have been able to directly or indirectly take advantage of the rising productivity and low wages of the developed world to dramatically lower costs and increase sales and profits. This poses two problems. First, many of these winners may not be manufacturing workers or even in manufacturing firms (for example, they may be providing R&D, design or advertising, or finance services to firms whose products are manufactured abroad).[3] Second, high-income individuals are hard to track through surveys, given their small number and non-response. Even if they do respond, capturing their gains is difficult given that much of their increase in wealth comes from stock options and investments rather than wages, and the super-rich also put considerable resources into concealing such gains. However, not clearly knowing who has gained from trade makes necessary redistribution that much different.

Given the paucity of research on the first issue, my comment will focus on the second: the effect of trade on price indices. Changes in the cost of living could potentially counteract the well-documented effects of import competition on nominal income inequality across regions. More generally, understanding the impact of trade on the denominator of real income is crucial to calculating the full effects of trade on inequality.

[1] I thank Penny Goldberg, Amit Khandelwal and Shinnosuke Kikuchi for useful discussions.

[2] See, for example, Autor, Dorn and Hanson (2013) or Pierce and Schott (2016).

[3] Many firms classified as manufacturers moved out of manufacturing in recent decades as their focus shifted to these high value-added services (see Ding et al., 2019, for the US). Magyari (2017) finds that employment in US manufacturing firms (rather than establishments) most heavily exposed to the China shock actually grew faster, driven by expansions in non-manufacturing tasks such as R&D, design, engineering and headquarters services, as well as shifts into less-exposed manufacturing industries. Also in the US, Bloom et al. (2019) find that highly educated coastal locations heavily exposed to the China shock saw total employment increase.

Effects of trade on price indices

Identifying the effects of trade on consumer price indices is challenging for two broad reasons: measurement and causal identification. In terms of measurement – and in contrast to objects such as wages and employment – many of the prices required to calculate welfare-consistent price indices are not observable. For many services, disentangling prices from quantities is extremely difficult. Even for expenditure categories with posted per-unit prices, improvements in quality lower the price index, as do increases in the variety available to consumers; yet neither of these changes are revealed from price quotes alone. Finally, the rich and poor shop at different stores, buy different brands and purchase different varieties. Thus, to understand differences in price indices across income groups, we require knowledge of the spending patterns of different income groups at a very granular level. The subsequent sections explore where these measurement issues generate knowledge gaps in our understanding of trade and inequality.

The second challenge is causal identification. The voluminous literature exploring the effects of greater import competition on regional wage and employment inequality exploits the enormous heterogeneity in industrial composition across commuting zones – motivating the use of shift-share instruments that interact predictors of national import penetration with initial industry composition.[4] In contrast, the similarity in consumption bundles and prices of tradables across regions means that the key effects of trade on price-index inequality are unlikely to play out across space. Thus, the literature to date has relied on product-level regressions where price-index changes for a particular product category are regressed on predictors of national import penetration. Jaravel and Sager (2019) pursue this approach to study the (relative) impacts of Chinese import competition across US consumer price index (CPI) product categories, as do Dorn and Levell (2021) using UK data. These papers find very substantial aggregate price effects with a 1 percentage point increase in Chinese import penetration lowering domestic prices by between 1% and 2%. Jaravel and Sager (2019) are further able to document that much of this price-index reduction is driven by reduced prices for domestic goods, with the authors highlighting mark-up reductions as the key mechanism.

These approaches rely very directly on shocks to import penetration being exogenous to existing price trends, at least conditional on controls.[5] This assumption is clearly violated for simple measures of Chinese import competition because, absent any controls for product-specific trends, the products groups that witnessed the largest growth in Chinese exports in the 2000s, such as electronics, also tended to be the products groups whose price indices fell most in each of the previous three decades, at least based on public release US CPI data that reports inflation by product category for around 200 product groups.[6] (Technology products in the UK analysis by Dorn and Levell (2021) appear to be similarly problematic and are excluded in their preferred specification.) Whether these identification concerns cloud estimates of the effects of trade on price-index inequality is less clear, and an issue I return to below.

[4] See, for example, Topalova (2010), Autor et al. (2013) and Kovak (2013).

[5] There is a lively debate on whether the regional analysis of import competition using shift-share instruments also implicitly requires that import shocks are exogenous (see, e.g. Goldsmith-Pinkham, Sorkin and Swift, 2020; Borusyak, Hull and Jaravel (2021). Borusyak et al. (2021) show that the standard regional regressions are equivalent to running a national product-level regression where variables are product-exposure-weighted national averages and the instrument is national shocks to import penetration by product. That said, endogeneity concerns are more severe and the bias of import competition shocks is likely to be larger when regressing product-level price indices on product-level imports.

[6] The raw CPI data by product group are available at https://www.bls.gov/cpi/data.htm but a concordance is required to go between product codes and the SIC codes used in Autor et al. (2013).

What we know about heterogeneity across income groups

For questions related to inequality, what matters are differences in inflation across income groups. At a minimum, we require expenditure data across product categories by income group. A priori, the fact that the poor tend to spend a larger share of their income on tradable goods[7] suggests they will benefit more from cheaper imports. However, working against this is the fact that within tradable categories, the rich spend relatively more on the most import intensive categories, such as electronics, and less on categories such as food that have small imported shares. Whether a trade reform favours the rich or the poor on the expenditure side is likely to depend on which goods and which trading partners see changes in import penetration. The analysis in Dorn and Levell (2021) uses expenditure shares from the UK 2001 Living Costs and Food Survey. The authors conclude that, at least for Chinese import competition, these two forces approximately cancel. Other analyses take account of imported intermediates embedded in domestic products and find a similar lack of inflation heterogeneity across income groups; see, for example, Breinlich et al. (2019), who investigate the effects of the depreciation of the UK pound after the Brexit referendum.

A shortcoming of analyses based only on CPI data and expenditure surveys is that CPI price quotes come from stores and no information on quantities is collected. Expenditure surveys do record expenditures and quantities but at too coarse a level to capture differences in the brands and varieties consumed by the rich and poor. If the rich disproportionately consume imported varieties within a product category, trade may still generate inequality via the price index. To fill this gap, 'home scanner' datasets collected by market research firms, such as GfK, Kantar and Nielsen, which ask consumers to scan purchases they bring home, have proven invaluable. Although primarily capturing supermarket retail (as other purchases are scanned less reliably by participants), these datasets contain both expenditures and quantities at the barcode level matched to basic household demographics including location and household income bin.

Several papers exploit granular data of this type. Levell, O'Connell and Smith (2017) do not find differences in the imported share of food consumption in UK home scanner data. Bai and Stumpner (2019) regress price indices for product categories covered by the Nielsen Homescan data in the US on Chinese import penetration (instrumented with European countries' Chinese import penetration similarly to Autor et al., 2013). They find that a 1% increase in the imported share of expenditure results in 0.17% lower inflation per year relative to a category experiencing no change in import penetration.[8] There is some evidence of heterogeneity across income groups. The effect size increases monotonically across income bins with the highest income bin experiencing a 20% larger reduction in inflation in import competing categories compared with the lowest bin. That said, the estimates are sufficiently imprecise that these differences are not significant. And recall that these indices primarily cover supermarket retail, a set of goods that is relatively less affected by Chinese import competition.

Returning to the discussion of confounding pre-trends, even if the main effect size is biased, it is not clear that the differences across income groups would be (i.e. the fact that for many decades prices have been falling faster in the sectors heavily exposed to Chinese imports does not imply that this pattern has been more pronounced for varieties purchased by the rich). Unfortunately,

[7] See, for example, Fajgelbaum and Khandelwal (2016).

[8] One-third of this reduction comes through greater variety, a channel I discuss in the section 'Harder-to-measure heterogeneity'.

the Nielsen sample made available to researchers only starts in 2004, precluding testing for these income-group-specific pre-trends.[9]

Borusyak and Jaravel (2018) provide perhaps the most comprehensive analysis of the effects of trade on price-index inequality. Essentially, they combine the two approaches described above, examining expenditure surveys and import intensities (from all countries, not just China) for US CPI products categories, but supplementing this analysis with more granular data for supermarket retail and passenger vehicles. Unlike Bai and Stumpner (2019) who regress Chinese import shocks on price indices, Borusyak and Jaravel (2018) identify imported varieties from all countries in the US Nielsen Homescan data (via a merge with the Economic Census and Customs data). They ask whether poorer consumers' bundles are more import-intensive, and similarly for passenger vehicles. Higher-income households (proxied by whether the household respondent has a college degree) do purchase more import-intensive varieties within product categories, at least for products in the Nielsen Homescan data and for passenger vehicles. They also spend a larger share of their goods expenditures on categories with high import shares. However, higher-income households' larger share of spending on services means that their expenditure is only slightly more tilted to imports than lower-income households (on average, 12.6% of expenditure is spent on imports, directly or indirectly, with a difference of 0.2% between college- and non-college-educated households). If all imports became cheaper, then differential effects across these household types would be less than 2% of the aggregate effect. Although these differences are about three times larger for Chinese imports than aggregate imports, they are still relatively minor.[10]

To conclude, even though there is still considerable uncertainty about the size of aggregate price-index effects due to increased import penetration, any distributional consequences due to the rich and poor consuming different categories of products, or due to different varieties within supermarket retail and passenger vehicles, are likely to be relatively limited. In the section 'Harder-to-measure heterogeneity', I discuss several less easily measured channels that could modify this conclusion.

That said, there are several mechanisms through which trade can alter price indices beyond changing the price and availability of imported goods. If the main impact of import competition comes from domestic firms reducing their mark-ups, as suggested by Jaravel and Sager (2019), then the import penetration methodology above may be missing more substantial impacts on inequality if these mark-up reductions are more pronounced for products disproportionately consumed by the rich.[11] And more generally, increased trade in a particular product category may induce price changes in other product categories, either through consumer substitution patterns or through changes in factor prices. Increased exports by domestic firms, beyond their very direct effects on wages, may also allow these firms to reap economies of scale, lowering domestic prices. Relatedly, a larger market size for exporters may raise the returns to innovation, providing consumers with new and/or cheaper products. The globalisation of capital markets has resulted in large flows of foreign capital flowing into US technology sectors, further spurring

[9] Broda and Romalis (2009) analyse Nielsen data from the period 1994–2005, but these data are not available through Nielsen's collaboration with the Kilts Center for Marketing at the Chicago Booth School of Business. Interestingly, they find that inflation rates were lower for the poor than the rich over this period, a pattern that reversed in the 2004–15 sample that spans the Great Recession (see, e.g. Jaravel, 2019; Argente and Lee, 2021).

[10] Poor households consume a higher share of Chinese imports within the Nielsen data but consume a lower share of categories that are intensive in Chinese imports (e.g. electronics).

[11] Jaravel and Sager (2019) suggest that the opposite may be true as price reductions appear to be larger for product categories that the poor consume disproportionately. However, the lack of quantity data accompanying the CPI price quotes makes it challenging to carry out a more complete distributional analysis using these data.

innovation.[12] How large such effects are, and whether they have differential effects for different income groups, is an open question.[13]

The finding that distributional effects are likely to be limited stands in some contrast to the evidence we have from household surveys regarding the impacts of trade on consumers in the developing world. Porto (2006) finds that price-index effects of Argentina's entry into the Mercosur regional trade agreement benefited the rich far more than the poor (although this was primarily driven by the specifics of Mercosur, which ended up raising food prices while lowering the prices of other tradables). In contrast, Marchand (2012) and Han et al. (2016) study the effects of trade reforms in India and China, respectively, finding that tariffs reduced prices for the poor more than for the rich. However, none of these studies allows for the rich to have higher import shares within broad product categories. Looking within product categories, Faber (2014) finds that the price index of the rich in Mexico declines more from NAFTA-induced reductions in import tariffs; however, in this case, the mechanism is through cheaper imported intermediates that disproportionately go into high-quality products consumed by wealthier households. Atkin, Faber and Gonzalez-Navarro (2018) explore the gains to the globalisation of the retail sector for Mexican consumers – specifically the entry of Walmart into Mexico – and find that consumption benefits are strongly concentrated on the rich who are much more likely to patronise foreign-owned supermarkets. These results are not inconsistent with the analysis above, given that, in the Mexican case, imports increased from countries at similar or greater levels of development. These countries are likely to have a comparative advantage in goods consumed by the rich. For example, Borusyak and Jaravel (2018) note that trade with the most developed countries does lead to substantially larger reductions in the price indices of the rich (with the difference in effects equal to about 16% of the aggregate reduction in the price index). Much of this is driven by their focus on passenger vehicles where the luxury segment is dominated by imported varieties from developed countries.

What we know about heterogeneity across regions

Much of the recent literature pertaining to trade and inequality has focused on regional inequality. For example, Autor et al. (2013) document large regional divergence in wages and employment as a result of Chinese import competition, and Dorn and Levell (2021) find something similar for the UK. An obvious question is whether countervailing changes in regional price indices may have moderated the impact on real income inequality across regions.

As many tradable goods have little regional price variation[14] and consumption bundles are broadly similar across regions in developed countries, there is limited scope for substantial regional responses in the price index of tradable goods. Consistent with this conjecture, Bai and Stumpner (2019) show that in the Nielsen sample, price-index responses to Chinese import competition were similar across four broad regions of the US. However, note this is not the same exercise as regressing a commuting-zone-specific price index on commuting-zone-import exposure that would be needed to directly answer the question posed in this section (at least for Nielsen products).

[12] See, for example, Akcigit et al. (2020).

[13] Acemoglu and Linn (2004) and Costinot et al. (2019) provide evidence that domestic market size affects innovation and prices, respectively, in the pharmaceutical industry. Jaravel (2019) shows that innovation in US retail categories is disproportionately directed to products purchased by wealthier consumers.

[14] See, for example, DellaVigna and Gentzkow (2019) for evidence of uniform pricing in US retail chains.

Even if tradable prices do not differentially respond across locations, households at different income levels may be better able to respond to wage reductions. Argente and Lee (2021) show that the Nielsen price indices of richer households grew less than those of poorer households during the Great Recession, as they were able to substitute out of expensive high-quality brands into cheaper substitutes and make more effective use of coupons and discounts. The poor were already consuming the more basic goods and taking advantage of discounts so had limited scope to adjust in this manner. This finding suggests that poorer households and poorer locations adversely affected by import competition may have seen larger declines in real income than richer households and richer locations affected by the same shocks.

What the above analyses do not touch upon are the prices of non-tradable goods, most notably housing, which is one of the largest expenditure categories. Given wage declines and unemployment in hard-hit locations, we would expect declines in rent (actual or imputed) as well as services that rely heavily on local labour inputs. Several papers explore housing price responses – changes in prices of other services are more difficult to measure, an issue I return to in the next section. Most directly, Feler and Senses (2017) regress housing prices and rents on the China shock shift-share instrument at the commuting-zone level and find significant reductions in housing prices and rents in heavily exposed locations relative to less-exposed locations.[15] For a $1,000 increase in Chinese imports per workers, house prices fell 5.4% and median rent fell by 2.5% relative to an unshocked location. This compares to a 2.14% decline in household income reported by Autor et al. (2013). The average consumer unit in the US spent 19% of expenditure on shelter (from the 1999 Consumer Expenditure Survey) and 65% of units own homes. So, roughly speaking, the reduction in the price of shelter is about 40% of the size of the reduction in nominal income. Of course, for someone already owning a home, the reduction in house prices also results in a negative wealth shock – a shock that is more pronounced for richer households that are less likely to rent.

Regional variation in the level of housing prices, rather than changes, also has implications for regional inequality generated by trade. Breinlich et al. (2019) note that London was relatively less affected by the devaluation following the Brexit referendum because the high share of income spent on expensive London housing costs meant that Londoners had lower expenditure shares spent on imports. More generally, 'superstar' cities such as London, New York and San Francisco have become increasingly attractive places for high-skilled service workers to agglomerate, with trade and globalisation surely playing a role. At least some of the regional (and aggregate) inequality that this has generated has been undone by soaring housing costs in these locations, as documented by Moretti (2013).[16]

Feler and Senses (2017) also document two more effects that have implications for regional price-index inequality. First, the China shock led to a significant increase in business closures, presumably the result of lower local demand. These closures reduced the variety of retailers and services available locally (an effect that may also be present within vendors[17]). As I discuss below, reduced variety raises the price index, exacerbating inequality. Pushing in the other direction, reduced local demand coupled with lower wages almost certainly lowered the prices of non-tradable goods and services that were still available.

[15] See also Xu, Ma and Feenstra (2019).

[16] As Moretti (2013) notes, it is also possible that rising housing prices reflect improvements in amenities in these cities. However, the fact that the relative wage of college-educated workers rose alongside their share of the population suggests that amenity improvements are not the primary driver.

[17] Handbury and Weinstein (2014) document substantial differences in the variety available within stores across small and large cities.

Second, in part due to declines in property and business taxes paid by richer households, spending on local public services declines in hard-hit communities. If these cuts affect service quality – and Feler and Senses (2017) show that both crime and student-to-teacher ratios increase – there are negative implications for real income, at least for poorer households who are net recipients of public-goods expenditures.

Harder-to-measure heterogeneity

The discussion in the two previous sections centred on elements of the price index that are straightforward to compute using price and expenditure data. Several studies drew on CPI data. Expenditure surveys provide expenditure shares on product categories by income group or region, while CPI price quotes reveal price changes for identical varieties within those categories. Putting the two together, it is straightforward to calculate standard price indices (e.g. a Laspeyres or Fisher ideal index). Home scanner data allow us to do the same for varieties within certain product categories. What these simple indices miss are changes in the quality of products and changes in the variety on offer (including changes in the number of vendors for a particular product or service). Many of the gains from trade on the consumer side come from access to a wider variety of goods or from being able to purchase cheaper goods of equivalent quality to domestic varieties. Thus, it is vital to properly account for variety and quality changes in order to capture these channels.

Statistical agencies such as the Bureau of Labor Statistics attempt to correct for quality changes for existing goods either by exploiting price differences between old and new varieties, if both are in the market at the same time (the 'overlap' method), or by using hedonic regressions for certain product categories. These quality changes are reflected in analyses that use CPI data, such as Jaravel and Sager (2019). The accuracy of these corrections varies greatly over product categories, with particular challenges measuring quality changes in service industries, such as healthcare, or durables experiencing rapid innovation on multiple dimensions, such as electronics. Novel products present an even more fundamental challenge and no correction is made for the consumer surplus generated by their arrival (although subsequent price changes do enter the CPI once price-quote collection starts, and some new products are partially accounted for if they 'replace' similar disappearing products for which price quotes were previously collected). Boskin et al. (1997), Moulton and Moses (1997) and Gordon (2006) provide a more detailed discussion of these CPI measurement issues.

In contrast, the market research companies collecting home scanner data do not provide quality corrections. Yet, because these datasets contain both price and quantity data, corrections for both quality and variety changes are possible given assumptions on demand. Feenstra (1994) shows that if preferences across varieties take the constant elasticity of substitution (CES) form, the change in the price index due to changes in product variety is parsimoniously revealed by the market share that the new products take, with the CES converting these market share changes to welfare changes. As many quality changes result in new barcodes being assigned, applying this correction to a barcode-level dataset, such as the Nielsen Homescan data, corrects for both quality and variety changes.[18] This approach is pursued by Bai and Stumpner (2019), as discussed above, with variety expansion accounting for one-third of the price-index reductions due to the

[18]Note that even if a barcode remains the same but the characteristics change slightly, such as the product becoming more durable, treating it as a new variety will produce the correct price-index change. The only limitation to this approach is that there must be at least one variety whose quality is known not be changing.

China shock in the Nielsen sample if they assume an elasticity of substitution of 5 across barcodes. Amiti et al. (2018) perform a related analysis using import data suggesting that two-thirds of the change in the US manufacturing price index due to China's entry into the World Trade Organization came from new Chinese varieties.

Although CES preferences are homothetic, if different income groups are assumed to have different (but fixed) preferences, income-group-specific price indices that accommodate variety gains can be calculated. Bai and Stumpner (2019) include such a variety term when documenting that price indices fell slightly more for higher-income households as a result of the China shock. (I am not aware of an analysis that considers the gains from variety across the income distribution due to trade with both developed and developing countries.) However, a major caveat is that the Nielsen sample and similar home scanner datasets provide poor coverage of big ticket durable items, such as televisions, appliances, computers and electronics. More generally, many expenditure categories, such as services and housing, are entirely absent from scanner datasets.[19]

Taking stock, CPI datasets cover the full consumption bundle but struggle to accurately measure quality changes for many categories and do not account for welfare gains from the arrival of novel products. Scanner datasets allow progress to be made on computing quality and variety corrections, subject to assumptions on demand, but omit the majority of expenditures.

Atkin et al. (2020) propose a methodology to recover changes in the full price index – accounting for price changes, variety gains, and quality improvements in all sectors – for households at every point in the income distribution, in contexts where detailed price and quantity data are not available for all product categories. Specifically, they examine budget shares spent on particular goods within a well-measured product category such as the food module covered by the Nielsen sample. The plot of these within-category shares against log income generates a 'relative' Engel curve. They first correct for shifts in these curves over time due to price changes within the food category where well-measured prices are available. Horizontal shifts in the resulting relative Engel curves reveal the change in the full price index for households at any point on the initial relative Engel curve, as long as preferences are quasi-separable – separable in the expenditure function rather than utility function. Applying this methodology to Indian inflation between 1987 and 2000, which was a period of major trade reforms, they find that CPI estimates focusing only on well-measured sectors miss the fact that inflation was much lower for rich households than for poor. This finding is consistent with the wealthy benefiting disproportionately from the arrival of new foreign varieties as a result of trade reforms as well as quality and variety improvements in service sectors. The fact that the gains disproportionately accrue to the rich resonates with Jaravel (2019) who makes the case that innovation in the US has been disproportionately directed at products consumed by the rich, and thus prices have risen more slowly for these households when accounting for their greater variety gains.

Product categories where these shortcomings are most serious

In analysing the price-index effects of trade, the shortcomings highlighted above are most serious for categories where we think trade has had pronounced effects on quality and variety. I now turn to discussing several such categories.

[19] Jaravel (2019) reports that purchases captured in the Nielsen data cover, on average, only 6.71% of household expenditure.

Technology-intensive goods, such as consumer electronics (including televisions and computers) and appliances, loom large. In terms of their quantitative importance, they are a significant share of expenditure, their import shares are high[20] and, given their high cost, they are consumed disproportionately by the rich. It is also a sector that has seen particularly rapid innovation, including the creation of entirely new products such as tablets, e-readers and streaming devices. Measuring quality improvements for upgrades to existing products is also extremely challenging. For these reasons, the Boskin commission's study of bias in the CPI estimated that existing CPI inflation estimates for these categories were likely the most overstated, with inflation as much as 5 percentage points lower in these sectors (Boskin et al., 1997).

For example, consider a new version of an iPhone. Given strategic pricing by Apple, the price change across versions is not very revealing of the magnitude of quality improvement using the 'overlap' method. Comparing crude characteristics such as the megapixels in the camera via hedonic regressions has only limited relevance to the quality of photographs the camera produces. Longevity is also important for durables of this type and this is difficult to measure early in a product's life. And genuinely new features, such as face unlock, cannot be accounted for through hedonics, or at least not at the time of the feature's introduction. There are also thorny issues related to the carefully curated customer perception of the product that may change over time through, for example, advertising and design. And of course, these challenges are all the more difficult when assessing the gains from the arrival of the first iPhone, a novel product. As briefly mentioned above and discussed in Jaravel (2019), these innovations have further implications for inequality as they may be targeted disproportionately to wealthy consumers purchasing more expensive varieties given the spending power and more inelastic demand (e.g. the iPhone Pro consumer rather than the iPhone SE consumer). Taken together, existing estimates of inflation in these product categories are likely overestimates, suggesting that overall inflation has been lower particularly for wealthier households.

One category of technology-intensive goods that has wider ramifications than simply affecting the price index is goods that are complementary with leisure, for example computers and games consoles. As discussed in Aguiar et al. (2021), reductions in quality-adjusted prices of these goods have shifted the leisure time of young men toward video gaming and recreational computing and, by raising the value of time, reduced their labour supply. Given that two of the three console producers are foreign, and almost all computer and console equipment is imported, trade surely has played a significant role.

Three other categories with substantial expenditures and large import shares are worth mentioning. Furniture, apparel and motor vehicles have import shares between 25% and 50% (Borusyak and Jaravel, 2018). Given the lower pace of innovation in this sector, measuring quality is easier and variety gains are likely to be smaller. That said, furniture and apparel are products where fashion and styles are important. Imports directly increase variety (e.g. the availability of Italian apparel brands raises welfare for US consumers, with the rich likely to gain disproportionately in this case). But offshore production also allows US and foreign designers to produce desirable clothes at much lower price points (e.g. 'Fast Fashion' that disproportionately benefits poorer consumers).

Services are a category that has garnered less attention from trade economists. However, trade in services constitutes almost 20% of imports in the US and 30% in the UK. Together, tourism and

[20]Borusyak and Jaravel (2018) calculate a 50% import share for computers and electronics in the US, including indirectly via intermediate inputs.

transportation services account for more than 40% of these imports in both countries. Given these numbers, falling costs of airline tickets generate non-trivial reductions in price indices that are picked up in CPI estimates. However, increases in the quality and variety of tourism experiences (e.g. exotic destinations becoming more accessible or existing destinations improving tourist infrastructure), as well as innovations such as AirBnB, Google Maps and Uber that substantially improve the quality of foreign travel experiences, also have the potential to substantially reduce price indices. These gains are unlikely to be distributed equally across income groups. On the one hand, international travel is a luxury, with the rich much more likely to travel abroad.[21] On the other hand, the rise of budget airlines in recent decades has favoured poorer cost-conscious travellers.

Relatedly, as globalisation leads more people to study abroad (education service imports) or seek medical care abroad (health service imports), the price index for these categories will decline. The former is likely to benefit rich households who are more likely to attend college, and the latter more likely to benefit poor households, at least in countries such as the US where medical treatment abroad is typically a way to access cheaper care.[22]

A final category of service imports that is increasing in importance is telecommunications, computer and information services. This category is particularly relevant for European consumers given US dominance in the websites and mobile phone applications used in Europe. That said, as of August 2020, the Chinese-owned short-video app TikTok had 100 million active US users – and a similar number in Europe – with the average user spending 52 minutes a day on the app.[23] Calculating the impact of websites and apps on the price index is fraught with difficulty, particularly as most do not charge consumers. However, the impacts on welfare are likely to be substantial and may vary by income group, given differences in time spent online by household types. Boik, Greenstein and Prince (2016) use data that track usage on 40,000 primary home consumers in 2008 and 2013 to document that US households that make less than $25,000 spent 92 more minutes online than those making more than $100,000. Relatedly, younger individuals (who tend to have lower incomes) spend almost three times as many hours a day on social media than older individuals.[24] Television and streaming services, such as Netflix, also increasingly expose viewers to imported content, with more-educated consumers likely to be more willing to watch such content.[25]

One final consideration is that trade and globalisation may change preferences themselves. This may occur through the process of habit formation as individuals consume more foreign products and thus habituate themselves to these varieties and brands (see Atkin, 2013), or through greater exposure to foreign products through travel. While welfare evaluation is conceptually difficult when preferences are changing, these forces are likely to further increase expenditure shares on imports. Nardotto and Sequeira (2021) also show that preferences for foreign varieties can decline if there is a backlash against globalisation. The Brexit referendum in the UK resulted in a

[21]For example, 69% and 58% of households earning less than $20,000 or $20,000–$35,000 a year, respectively, reported never travelling outside the US. The equivalent percentage was 25% for households earning $100,000 or more (see https://www.statista.com/statistics/668013/international-vacation-frequency-of-us-adults-by-income/).

[22]150,000–320,000 US travellers list health care as the reason for travelling abroad each year, with Latin American and Caribbean destinations accounting for 65% of travellers (Chambers, 2015).

[23]See https://wallaroomedia.com/blog/social-media/tiktok-statistics/.

[24]See https://www.digitalinformationworld.com/2019/01/how-much-time-do-people-spend-social-media-infographic.html.

[25]In 2020, US viewership of foreign language shows increased by 50% (https://www.bloomberg.com/news/articles/2020-12-10/netflix-s-foreign-language-shows-see-popularity-soar-in-the-u-s).

sudden 13% decline in the market share of European Union products relative to similar UK products marked with a Union Jack or featuring a British origin in the product name (despite no change in relative prices or trading conditions at the time and conditioning on shopper fixed effects). Beyond potentially direct effects on welfare, these preference changes change consumers' exposure to future trade shocks.

References

Acemoglu, D. and Linn, J. (2004), 'Market Size in Innovation: Theory and Evidence from the Pharmaceutical Industry', Quarterly Journal of Economics, 119, 1049–90.

Aguiar, M., Bils, M., Charles, K. K. and Hurst, E. (2021), 'Leisure Luxuries and the Labor Supply of Young Men', Journal of Political Economy, 129, 337–82.

Akcigit, U., Ates, S. T., Lerner, J., Townsend, R. R. and Zhestkova, Y. (2020), 'Fencing Off Silicon Valley: Cross-Border Venture Capital and Technology Spillovers', National Bureau of Economic Research Working Paper No. 27828.

Amiti, M., Dai, M., Feenstra, R. C. and Romalis, J. (2018), 'How Did China's WTO Entry Benefit US Prices?', National Bureau of Economic Research Working Paper No. 23487.

Argente, D. and Lee, M. (2021), 'Cost of Living Inequality During the Great Recession', Journal of the European Economic Association, 19, 913–52.

Atkin, D. (2013), 'Trade, Tastes, and Nutrition in India', American Economic Review, 103 (5), 1629–63.

Atkin, D., Faber, B., Fally, T. and Gonzalez-Navarro, M. (2020), 'Measuring Welfare and Inequality with Incomplete Price Information', National Bureau of Economic Research Working Paper No. 26890.

Atkin, D., Faber, B. and Gonzalez-Navarro, M. (2018), 'Retail Globalization and Household Welfare: Evidence from Mexico', Journal of Political Economy, 126, 1–73.

Autor, D. H., Dorn, D. and Hanson, G. H. (2013), 'The China Syndrome: Local Labor Market Effects of Import Competition in the United States', American Economic Review, 103 (6), 2121–68.

Bai, L. and Stumpner, S. (2019), 'Estimating US Consumer Gains from Chinese Imports', American Economic Review: Insights, 1, 209–24.

Bloom, N., Handley, K., Kurman, A. and Luck, P. (2019), 'The Impact of Chinese Trade on US Employment: The Good, the Bad, and the Debatable', Working Paper.

Boik, A., Greenstein, S. and Prince, J. (2016), 'The Empirical Economics of Online Attention', National Bureau of Economic Research Working Paper No. 22427.

Borusyak, K., Hull, P. and Jaravel, X. (2021), 'Quasi-Experimental Shift-Share Research Designs', Review of Economic Studies, in press, https://doi.org/10.1093/restud/rdab030.

Borusyak, K. and Jaravel, X. (2018), 'The Distributional Effects of Trade: Theory and Evidence from the United States', Harvard University working paper.

Boskin, M. J., Dulberger, E. R., Gordon, R. J., Griliches, Z. and Jorgenson, D. W. (1997), 'The CPI Commission: Findings and Recommendations', American Economic Review, 87 (2), 78–83.

Breinlich, H., Leromain, E., Novy, D. and Sampson, T. (2019), 'Exchange Rates and Consumer Prices: Evidence from Brexit', Centre for Economic Performance (CEP) Discussion Paper No. 1667.

Broda, C. and Romalis, J. (2009), 'The Welfare Implications of Rising Price Dispersion', University of Chicago, mimeo.

Chambers, A. (2015), Trends in US Health Travel Services Trade, Washington, DC: US International Trade Commission.

Costinot, A., Donaldson, D., Kyle, M. and Williams, H. (2019), 'The More We Die, The More We Sell? A Simple Test of the Home-Market Effect', Quarterly Journal of Economics, 134, 843–94.

DellaVigna, S. and Gentzkow, M. (2019), 'Uniform Pricing in U.S. Retail Chains', Quarterly Journal of Economics, 134, 2011–84.

Ding, X., Fort, T. C., Redding, S. J. and Schott, P. K. (2019), 'Structural Change Within versus Across Firms: Evidence from the United States', Working Paper.

Dorn, D. and Levell, P. (2021), 'Trade and Inequality in Europe and the US', IFS Deaton Review of Inequalities.

Faber, B. (2014), 'Trade Liberalization, the Price of Quality, and Inequality: Evidence from Mexican Store Prices', UC-Berkeley Working Paper.

Fajgelbaum, P. and Khandelwal, A. (2016), 'Measuring the Unequal Gains from Trade', Quarterly Journal of Economics, 131, 1113–80.

Feenstra, R. C. (1994), 'New Product Varieties and the Measurement of International Prices', American Economic Review, 84 (1), 155–57.

Feler, L. and Senses, M. Z. (2017), 'Trade Shocks and the Provision of Local Public Goods', American Economic Journal: Economic Policy, 9, 101–43.

Goldsmith-Pinkham, P., Sorkin, I. and Swift, H. (2020), 'Bartik Instruments: What, When, Why, and How', American Economic Review, 110 (8), 2586–624.

Gordon, R. J. (2006), 'The Boskin Commission Report: A Retrospective One Decade Later', National Bureau of Economic Research Working Paper No. 12311.

Han, J., Liu, R., Marchand, B. U. and Zhang, J. (2016), 'Market Structure, Imperfect Tariff Pass-Through, and Household Welfare in Urban China', Journal of International Economics, 100, 220–32.

Handbury, J. and Weinstein, D. E. (2014), 'Goods Prices and Availability in Cities', Review of Economic Studies, 82, 258–96.

Jaravel, X. (2019), 'The Unequal Gains from Product Innovations: Evidence from the US Retail Sector', Quarterly Journal of Economics, 134, 715–83.

Jaravel, X. and Sager, E. (2019), 'What are the Price Effects of Trade? Evidence from the US and Implications for Quantitative Trade Models', Centre for Economic Performance (CEP) Discussion Paper No. 1642.

Kovak, B. K. (2013), 'Regional Effects of Trade Reform: What Is the Correct Measure of Liberalization?', American Economic Review, 103 (5), 1960–76.

Levell, P., O'Connell, M. and Smith, K. (2017), 'The Exposure of Households' Food Spending to Tariff Changes and Exchange Rate Movements', Institute for Fiscal Studies (IFS) Briefing Note BN213, https://ifs.org.uk/publications/9563.

Magyari, I. (2017), 'Firm Reorganization, Chinese Imports, and US Manufacturing Employment', Columbia University.

Marchand, B. U. (2012), 'Tariff Pass-Through and the Distributional Effects of Trade Liberalization', Journal of Development Economics, 99, 265–81.

Moretti, E. (2013), 'Real Wage Inequality', American Economic Journal: Applied Economics, 5, 65–103.

Moulton, B. R. and Moses, K. E. (1997), 'Addressing the Quality Change Issue in the Consumer Price Index', Brookings Papers on Economic Activity, 1:1997, 305–66.

Nardotto, M. and Sequeira, S. (2021), 'Identity, Media and Consumer Behavior', Centre for Economic Policy Research (CEPR) Discussion Paper No. DP15765.

Pierce, J. R. and Schott, P. K. (2016), 'The Surprisingly Swift Decline of US Manufacturing Employment', American Economic Review, 106 (7), 1632–62.

Porto, G. G. (2006), 'Using Survey Data to Assess the Distributional Effects of Trade Policy', Journal of International Economics, 70, 140–160.

Topalova, P. (2010), 'Factor Immobility and Regional Impacts of Trade Liberalization: Evidence on Poverty from India', American Economic Journal: Applied Economics, 2, 1–41.

Xu, Y., Ma, H. and Feenstra, R. C. (2019), 'Magnification of the "China Shock" Through the U.S. Housing Market', National Bureau of Economic Research Working Paper No. 26432.