AER Referee Reports (submitted December, 2014)

Note: My responses to comments are in blue. Also, the page numbers correspond to the previous draft which I submitted to the AER, so may have changed. I attach these so that if I happen to get the same referees again, they can see how I changed the paper based on their comments, and also since it might be helpful for referees to see the issues raised by others.

Referee #1:

Referee Report on "Relative Prices, Hysteresis, and the Decline of American Manufacturing"
This paper assesses and quantifies the role of real exchange rates on the decline of American manufacturing. It establishes that real exchange rate movements had a larger effect in employment, output and other variables on manufacturing sectors that were more exposed to trade. The author concludes that "the appreciation of US relative unit labor costs can plausibly explain more than two-thirds of the decline in manufacturing employment in the early 2000s".

Response: This summary is mostly fine, but misses what I see as the chief contribution of the paper -- hysteresis. As stated in the paper, "The chief contribution of this paper lies in documenting the phenomenon of hysteresis at a disaggregated level for 437 SIC manufacturing sectors for both the 1980s and the 2000s." I see this as the chief contribution of the paper because the implicit model most trade economists use to think about trade, and trade policy issues such as China's currency undervaluation in the 1990s and 2000s, does not incorporate the potential for hysteresis. It's a result, which, if true, has implications beyond the field of trade.

Comments
Overall, I found hard to interpret the results in the paper. Here are a few reasons. (1) Different sectors are exposed to different exchange rates because they are exposed to demand in different countries. Given the identification strategy depends on how exchange rates interact with how open sectors are, I kept asking myself why the author focuses on an economy wide exchange rate rather than sector specific exchange rates. I do not mean using sector-specific exchange rates is the way to go – I also have questions about the conditions in which the use of sector-specific exchange rates is the right way to go – but I am definitely puzzled about why is the interaction between economy wide exchange rates and the degree of openness of the industry the right thing to be looking at, given the question. Can the author write a model to make precise under what conditions economy wide RER interacted with openness of the sector is the right approach to answer the questions he is after?

Response: This comment is fair enough, and something I discussed at length in footnote 34 on page 24 (footnote 41 in the revised version) and a question I struggled with myself when writing the paper. Essentially, I didn't get much more mileage out of the sector specific rates and so I didn't feature them as prominently. Note that neither Revenga (1992) or myself were able to create sector-specific prices (wages relative to productivity for the RULC index, or sector specific output prices for both the US and trading partners in her case) – Ravenga used *aggregate* Producer Price Indices (Revenga 1992, p. 282), while only the weights were sector specific. Thus the sector-specific exchange rates this referee suggests using are also clearly a sub-optimal measure. Basically, I feared that sectors with apparently very high WARULC's could have much lower-than-average RULCs themselves, and in that case the higher WARULC's would be partially spurious. We have data on US sector specific RULCs, but not the sector specific WARULCs of US trading partners, and so if we include sector specific-wages and productivity in the sector-specific exchange rates, then you'd end up with a biased measure. This might be why I find that the action mostly comes from movements in the economy-wide rate, although it remains a bit of a puzzle. This finding doesn't necessarily contradict Revenga. While she found that sectoral import prices were

correlated with employment declines from 1981-1985, these two variables could be correlated for demand-side reasons rather than due to actual movements in exchange rates. Unless I missed it, Revenga did not report the direct relationship between sectoral exchange rates and employment over this period after controlling for sectoral import shares.

(2) Appreciations hurt import-competing sectors lowering the price of imports, but it can also help import competing sectors that import a large fraction of its inputs. This is not taken into account in the present paper. Instead, only an openness variable (that is never defined in the paper) is used. I assume that the author means that sector-specific openness is equal to sectoral imports plus exports over production, or a similar measure.

Response: Note that in Table II (now Table I) there is a control for imported intermediate inputs interacted with RER movements, and that I discuss this in the text on page 21 (now page 20). Also, in footnote 30 I did say that openness is an average of import penetration and export share, although I did not define import penetration. That definition has been restored – thank you for the catch.

(3) If openness bundles exports and imports, then I do not see why one would not want to separately treat import competing from export oriented industries.

Response: Previously, in the appendix (page 54) I did them separately, now I have moved this table to the main text (Table II). Theoretically, the impact on imports and exports should be symmetric, which is what I found for the entire sample.

Table II has several specifications interacting the level of WARULC (author's measure for RER), with openness. The LHS variable, though, is the change in sector-level employment. How should we interpret this coefficient?

Response: See the bottom paragraph of page 21 (now page 18).

Why would the level of exchange rates affect changes in employment?

Response: See page 6: "The reason is conceptually easy to understand -- if unit labor costs were the same in the US and China, there would be no economic reason to move production, particularly as this could entail substantial fixed costs. On the other hand, when US unit labor costs are 50% higher than in trading partners, there is clearly an economic incentive to shift production, while firms already located abroad would have a competitive advantage." Note that the vast majority of economic models, including the one in the appendix, imply the relative prices matter and not simply the change in relative prices.

Table II should control for trends in sector-specific employment. The author can include a few lags of sector-specific growth in employment.

Response: Note that the dependent variable in Table II (note, this is now Table 1) is the log change in manufacturing employment, and thus the sectoral fixed effects control for trends in sector-specific employment. Perhaps the referee meant trends in sector-specific employment growth (the second derivative), which I have now added in the Appendix as a robustness check (along with several lags of the dependent variable), though it is not clear to me that these are intuitive controls. Note that I have plotted pre-and post-treatment trends in Figures 3, 4, 7, 8, 13, and 14, and they clearly show that time trends (or even second derivatives in trends) do not drive the results.

The notes in Table II read: "All regressions are weighted by initial sectoral value added". I have two observations about this. First, the weighting seems arbitrary. What is the reason why the author wishes to weight for value added? Since the author does not show the unweighted regressions, it gives the impression that the unweighted regressions must have given quite different results. To dilute this concern, the author should show unweighted regressions in addition to weighted regressions (if the author justifies the need for weighting). Regarding the need for weighting: I can only think of heteroskedasticity as a reason for weighting. In that case, weighting would mean doing GLS or feasible GLS in order to obtain more efficient estimators. I do not see why value-added would help the author to deal with heteroskedasticity here.

Response: Several points here: (1) In the paper I stated "The results do not appear to be sensitive to the choice of weights, as qualitatively similar results attain when weighting by average value-added, employment, or shipments, although the key coefficient is the largest when weighting by employment or when not weighting." (2) since the sectoral classifications are somewhat arbitrary, why on Earth would we not want to weight by the size of sectors? Heteroskedasticity is not the only reason to do weighted regressions. For example, Autor, Dorn, and Hanson (2013) report that their "Regression models are weighted by the start of period CZ [commuting zone] share of national population", because they'd naturally want to assign more weight to larger commuting zones just as I'd want to place more weight on larger sectors (which may be comprised of more sub-sectors). (3) Lastly, as an empirical economist, I typically find that weights do not matter that much – this seems like at most a minor issue. In any case, I have now included this as a robustness check in the main text of the paper (Table II).

Revenga (1992) has a paper studying the effect of import competition on American manufacturing employment and wages. She directly uses changes in sector-specific prices as measures for changes in import competition, and instruments these price changes with sector-specific exchange rate changes. I would like to have seen the current paper making a connection with Revenga (1992).

Response: OK, this is fair enough, I will cite Revenga since I also create sector-specific indices that she was one of the first to compute. There are many papers about the exchange rate shock in the 1980s and I do not have space to cite them all, although none of them, including Revenga, document hysteresis, and note that I did not find any papers on the appreciation in the US dollar in the late 1990s and early 2000s. Note that I did cite the Klein, Schuh, and Triest (2002) overview of the literature to that point, which includes Revenga (1992).

Importantly, the results in the paper are based on difference-in-differences methodologies. As such, the paper can only estimate relative effects of RER. The methodology relates how changes in RER interacted with sectoral openness explain changes in sectoral employment around the trend. RER has general equilibrium effects that are absorbed in the time fixed-effects. This is a well-known limitation of D-in-D.

Response: The assumption I make is that non-traded manufacturing sectors should not be impacted in a large way, particularly after controlling for demand growth, which controls for business cycle effects. In any case, the direct employment impacts I get from RER movements in the early 2000s are already more than twice as large as Autor/Dorn/Hanson, and this comment would imply that I'm perhaps understating the impact, which would only make my results more salient. However, I do believe that non-traded sectors were mostly unaffected by RER movements.

Overall, the writing could me more fluid, and the introduction shorter and more to the point. The empirical approach should also be more rigorous, and preferably connected to a (simple) model of trade.

Response: I have such a model in the appendix. It is in the appendix because several referees at the QJE wrote that the impact of RER movements is intuitive. I'd have to say I agree. Revenga (1992) does

assume a simple model, but her estimating equation is assumed not derived, whereas my estimating equation is implied by the model in the appendix. Asking for a model (or asking to remove a model from a paper) seems to be a default comment referees make of empirical work when they need to fill up a referee report.

PS: I am writing this report on May 18, and feel I should apologize with the author and editor for taking so long in delivering it. I had too many obligations over the semester and a continuously growing backlog of papers to referee. My apologies.

Overall Comments on Referee Report #1: I don't mind the response time, but I did note that six out of the first seven comments this referee made (out of 10 total) appear to have in fact been addressed directly in the text of the paper submitted. Also, while I gathered that this referee had a negative view of the paper, I did not really get a sense why, as most of the issues raised were relatively simple robustness checks (and checks that mostly were already in the paper submitted). The focus of the report on regression weights, for example, seemed bizarre. Nevertheless, I was able to use this report to improve the paper, and for that I am grateful.

Referee #2:

Review of "Relative Prices, Hysteresis, and the Decline in American Manufacturing"
This paper studies the role of real exchange rate (RER) appreciation in the collapse of U.S. manufacturing employment in the early 2000s. The paper uses new measures of RER that the author develops in a companion paper (Campbell 2014) and compares them to existing measures. The main identification strategy uses a difference-in-differences approach based on variation in RER for the U.S. over 1973-2009 and variation in openness to international trade across U.S. manufacturing sectors in the cross-section and over time. The author complements this strategy with an international perspective, comparing the outcomes for the U.S. to Canada and other OECD countries in a "triple-differences" approach and exploiting a quasi-experimental setting in Japan. The main findings document a negative impact of RER appreciation on manufacturing employment, which can account for two-thirds of the decline in U.S. manufacturing employment in the early 2000s. This effect is persistent and illustrates hysteresis.

Measurement
The main premise of the paper is the superiority of using a "Weighted Average Relative Unit Labor Cost" (WARULC) index instead of existing measures of RER. The main problem for existing measures of RER, both relative unit labor cost (RULC) and CPI-based measures, is compositional change in the importance of trading partners over time, which makes existing fixed trade weights outdated, and missing data on developing countries. Hence, existing measures of RER understate the dollar appreciation in the early 2000s and thus underestimate the relationship between RER and manufacturing employment. The companion paper Campbell (2014) provides a detailed discussion of the new measurements, but what is missing in the present paper is a brief discussion of some caveats, for example in terms of using weights based on current period trade shares instead of fixed trade weights. Trade shares are endogenous to unit labor costs and this might for example overstate the dollar appreciation in the early 2000s when China's trade share increased dramatically.

Response: I do not see the caveats arising from using an RER index with variable trade weights as being all that major. For example, if China is undervalued, the US export share should fall (or stay depressed relative to China's GDP growth), and thus it isn't necessarily the case that having an undervalued RULC with one trading partner would lead to an increase in their trade weights, although it might. In any case, if China is undervalued and their trade weight increases, the WARULC index would show a decline in competitiveness, just as it should. I don't see why one would prefer an index with fixed weights here,

which will not show any decline in competitiveness due to the rise of China or an increase in trade with an undervalued country. Additionally, I always use lagged values of the RER, so the current period trade weights never show up in any of my regressions. It is worth noting that almost all RERs use time-varying trade weights, including the OECD, the BIS, and the Federal Reserve, and other recent contributors such as Fahle, Thomas, and Marquez (2008) who have computed their own indices.

One potential fix here is that now I have computed a GDP-Weighted Average Relative Price Index, which should alleviate concerns about endogeneity in the weights in bilateral trading relationships. A second related fix is that I have also added a new robustness check using the Fed's RER index interacted with Import Penetration ex-China, and then interacted Chinese Import Penetration with a bilateral RER measure separately. While the Fed's index does use variable weights, the indexing method assures that China's rising trade share itself won't influence the level of the index. This measure is also almost identical to the IMF's CPI-based RER measure which has fixed trade weights, and still leaves the basic conclusion unchanged.

The IMF's RULC index uses fixed trade weights instead. If this is not an issue in practice, it is also worth pointing out. The details are important here because the author draws conclusions about the explanatory power of RER for the decline in manufacturing employment. Maybe a more conservative bound would be useful.

Response: Yes, and the IMF's weights do not include China, and give Japan a 20% weight, since they were based on outdated trade weights from the 1990s. As mentioned above, almost every other organization that creates RER indices use time-varying trade-weights, and most papers which study the impact of RER movements on trade or manufacturing employment also use time-varying weights. For example, Klein, Schuh, and Triest (2002), the closest paper in the literature to this one, uses the Fed's index which also has time-varying trade weights.

Moreover, the introduction of these measures is difficult to understand. The author uses abbreviations and terminology without sufficient explanation (ICP benchmarks, Divisia referring to the Fed's index in Figure 2, weights in WARULC footnote 23 unclear). The differences between the new WARULC measure and existing RULC measures are essential to the paper and should be clear from the beginning, despite the companion paper.

Response: Fair enough – I have taken more care to define terms in the updated draft.

A second important point is the measurement of manufacturing employment, which is based on sectoral data from the BEA. The main caveat is that there is a large set of plants that are engaged in manufacturing-related activities but are considered part of the wholesale sector and firms have been reclassified from manufacturing into wholesale and services. Bernard and Fort (2013) argue that employment at "factoryless good producing firms" has increased over time and hence the aggregate decline in manufacturing employment that is used in the current paper might be overstated.

Response: Interesting point. Three points though. First, even if these were truly manufacturing firms, one still needs to explain the decline in the traditionally-measured manufacturing sector. Second, even their largest estimates of the jobs gained in these factoryless good producing firms would shrink the number of overall manufacturing jobs lost in the early 2000s from 3 million to 2.4 million, which is still a large number of jobs lost in a short amount of time. Third, the US economy, particularly in terms of employment, did terribly during the 2000 to 2007 period despite there being a housing bubble and despite

historically low interest rates during this period. This fact is not consistent with a story in which the collapse in manufacturing was all about mislabeling manufacturing industries as services.

This argument is related to the process of deindustrialization and reclassification of firms as observed in the U.S. (Pierce and Schott 2012) and in other countries (Bernard, Smeets and Warzynski 2013). It would be useful if the author could mention this recent debate and relate the results to these more robust measures of manufacturing employment or use employment for a stable manufacturing sample over time based on the Longitudinal Business Database.

Response: Pierce and Schott (2015) imply that the reclassification of firms from NAICs to SIC had nothing to do with the observed decline in manufacturing employment, as they create consistent samples (as do I). In the main tables of the paper I excluded publishing-related sectors for the entire period due to the fact it that they were later reclassified out of manufacturing, and in the appendix I show that the results are little-changed when these sectors are added back in (however, there was also a typo in Table IX – now Table II -- below "Defense Related Sectors" should read "Publishing" and not "Full Controls"). Bernard, Smeets and Warzynski (2013) is a very interesting paper in which they note that, for Denmark, many manufacturing firms began to switch to services, and that by the end, the aggregate accumulation amounted to nearly 10% of manufacturing workers. One issue here is that, in some years, just as many firms switched in to manufacturing as those who switched out, point being that firms switch industries all the time, and so doing one-sided counting can be misleading. It was also the case that in the period before manufacturing declined, that there were plenty of out-switchers, and if you added back in, then manufacturing employment would have previously been increasing instead of flat. Thus, this seems to explain none of the change in trends in manufacturing employment, but rather would be a factor which would raise the trend in manufacturing employment in all periods. Lastly, the paper lacks evidence that firms who switched out of manufacturing were actually doing manufacturing-like activities.

In general, the approach and the findings of the paper are very interesting and compelling. Yet I thought the paper was not always clear in terms of labeling, notation, and definitions. As another example, "openness" is a key measure of the analysis and should be defined in the very beginning (instead of indirectly in footnote 30); "relative openness" is introduced in section 3 but the benchmark is unclear. I list more detailed examples in the last section.

Response: I'm glad to hear that the results were interesting and compelling. I apparently deleted the full definition of openness which contained the definition of import penetration during my last revision – I will certainly fix.

Identification strategy

Identification in the paper is primarily based on differences in openness across sectors and over time. The author discusses extensions such as using additional variation in the share of labor costs in value added or using import versus export weighted RER measures in the appendix. Moreover, international comparisons and policy variation is used to provide additional evidence for the main mechanism.

In general, one might be concerned about mean reversion when using openness lagged by one period as the main measure of exposure. A positive global demand shock might lead to higher exports and a temporary increase in employment for example. As demand decreases again, the sector with higher previous openness mechanically experiences a stronger reduction in employment. Controlling for demand change and TFP change does not fully account for the level of these shocks in time t-1. This is a concern with specifications in growth rates that are used in all main regressions. The paper and online appendices provide a variety of robustness and falsification exercises, so using longer lags to measure exposure should be added to those.

Response: This is a fair point. But note that Figure's 3 and 4 use fixed categories of openness (Figure 3 uses 1972, and Figure 4 uses openness in 1989), and thus the results do appear to be robust to this concern. As a robustness check, I tried relative openness at 2, 3, 4, and 5 year lags, and the results hold. I also used relative openness in 1972 for the panel regression in the full period, and still got statistically significant results, although the coefficient shrank considerably.

Results

The main findings in Table II illustrate the robust mechanism between RER appreciations and manufacturing employment. The results on various manufacturing outcomes support this evidence but it is surprising why TFP should decrease more for more open sectors. In theory, competitive pressure, offshoring and employment reduction could lead to an increase in productivity. Moreover, it would be nice to complement the decline in investment with an increase in FDI as firms are more likely to relocate production to low wage countries.

Response: Yes, I agree that the results on TFP are really quite striking. If manufacturing firms have overhead costs, then a decline in sales should lead to declines in TFP, so I wouldn't call these results counterintuitive. I can explain this in more detail in the next draft.

For the results on job creation and job destruction, the results seem somewhat different from the findings in Faberman (2008), who argues that the 2001 recession was followed by a persistent decrease in the job creation rate but a quick recovery of the destruction rate after 2001, so maybe the present results could be put into perspective.

Response: This would be a good point to address in my next draft. I think this could be explained by the elevated RER through the mid-2000s, and perhaps also to input-output effects, as manufacturing demand never came fully back in the 2000s, since industries that use other manufactures had already relocated.

The international comparison is a nice way of providing robustness checks but the regression specifications and the results need more explanation and interpretation. The results in Table V (1) suggest 25% larger difference in employment growth for a sector with twice the level of openness than the average sector, these results need further interpretation, for example by referring back to the actual changes in sectoral employment in different countries.

Response: Yes, I can do this.

In terms of specifications, it is unclear why the triple–difference estimation uses the level of WARULC and subtracts 1 (or 0.85 for Canada) instead of being consistent across sections and using log(WARULC) to make the results for Canada comparable to Table II for example.

Response: Good point. I did both sets of regressions both ways, it isn't significant. (Taking the log of a variable is *approximately* the same as subtracting one, of course, according to a Taylor expansion.) I did it this way in this case because, in general, country-specific factors (such as tariffs) mean that the equilibrium RER value needn't be one. Thus, I'd use behavioral rules to judge the equilibrium. In Canada's case, it seems this value is about .85. The results don't change if you use .9 or .8 (or .95) instead. I've added to the explanation of this in the paper.

Table VI for Japan could be extended to show the same mechanism as before by interacting the Post-Plaza Accord dummy with previous exposure to Japanese imports.

Response: Note that the dependent variable is not employment in this table, but import penetration. I'd be wary of judging the impact of Japan on employment, because Japanese imports were heavily concentrated in sectors with lots of imports from elsewhere.

The last section argues that the stable Japanese import penetration in the U.S. after the Yen appreciation is "validation" (p.34) for hysteresis. Yet there are other prominent explanations for this finding in the literature, in particular intra-industry trade and quality differentiation. A firm-level analysis of trade flows would be necessary to conclude that there is persistence in location decisions of firms; the aggregate evidence is only suggestive.

Response: I agree that it doesn't tell us the mechanism of hysteresis/persistence (is it location decisions? or market-specific investments? Market-specific learning-by-doing?), but I would stand by the assertion that it does indicate some kind of persistence. Aggregate variables can exhibit persistence just as much as firm-level observations can. However, I may not be understanding this comment fully.

In terms of the online appendix, the main point of the model is that current trade levels depend on the history of entry (and exit) conditions because of sunk costs. This implication relates to the focus on persistent effects of temporary RER shocks in the empirical analysis. Yet the model has many testable implications at the firm level about timing of entry, and changes in exports at the extensive and intensive margin that cannot be tested with the data available in this paper. The appendix does not explain why this particular framework with heterogeneous firms is chosen and which assumptions are needed to extend it to compare labor demand across sectors in general equilibrium as suggested by the empirical analysis. Given the fact that the recent literature has emphasized offshoring activities and global supply chains, the path dependence of these sourcing decisions could be modeled instead.

Response: Fair enough. The point of the model is simply to show that sunk costs lead to persistent impacts of RER shocks. The Melitz model is widely used, and I can get an analytical solution using it, including for a "dynamic gravity equation". While I agree it's possible to write down a model with persistence in sourcing global supply chains, I do wonder how much value this would add to this paper.

Some more detailed comments while reading the paper:
- Figure 3: How different are the top 25 and bottom 50 percent of sectors in terms of initial openness in 1972 and 1985? What role do intermediate inputs play here, what if firms are grouped by share of imported intermediates instead?

Response: I have now reported the cutoffs in the notes under the figures. The answer to the second question is that you don't see much action. Controlling for intermediate imports separately doesn't change the picture, however.

- Figure 4: It seems that the timing is interesting as well. The decrease in employment starts as WARULC increases above 1.2. Still sectors are much more open in the 2000s, so the question is to what extent they can further reduce employment.
Response: This is true. The sectors are more open in the 2000s, but the employment in those sectors is a bit less.
- Footnote 23: The weights omega should be defined.
Response: Fixed.
- Table I: The year 1972 is missing but relevant for Figure 3 for example.
Response: There was no data on duties or the cost of insurance and freight until 1974, so I began with that instead. Seemingly not a big difference.
- Figure 6(b): implies that import competition is the driving force. Why use openness throughout?

Response: I found that it isn't necessarily the case that imports are the driving force. When the RER appreciates, in the 1980s, exports actually decreased, and in the early 2000s, they were flat even though they had been increasing (just as overall shipments were increasing). Overall, I don't find a statistically different result for either import penetration or the export share, although the impact on imports is slightly larger (this regression was formerly in the appendix, but I moved to the main part of the paper).
- P.19 cites Pierce and Schott (2014) and Pierce and Schott (2013), the references list Pierce and Schott (2012).
Response: Fixed.
- P.23: to [do] fairly poorly
Response: Fixed.
- Table IV does not show shipments as mentioned in the text. Instead it shows TFP, which the text refers to the appendix.
Response: Fixed.
- Figure 10(a) y-axis label? How do we see the 2.07 million here?
- Figure 11 (c) and (d) legend 1979-1986 wrong
Response: Fixed.
- P.33: Just as China [as] has become
Response: Fixed.
- P.34: relative to the period when the Yen was undervalued
Response: Fixed.
- Figure 13: same axes for US and UK graphs
Response: Fixed.
- P.47: why not use overhead costs of exporting as well? Then there would be exit from exporting instead of restricting downwards adjustment at the extensive margin of trade to firm exit.
Response: Parsimony. They are not needed to generate persistence or to create a dynamic gravity equation, and including them would not change the basic story.
- P.49: Using p notation for both prices and share of surviving firms is very confusing
Response: Agreed, fixed.
References
Bernard, Andrew B., and Teresa C. Fort. "Factoryless goods producers in the US," No. w19396. National Bureau of Economic Research, 2013.
Bernard, Andrew B., Valerie Smeets and Frederic Warzynski, 2014. "Rethinking Deindustrialization," Economics Working Papers 2014-14, School of Economics and Management, University of Aarhus.
Douglas L. Campbell, 2014. "Through the Looking Glass: A WARPed View of Real Exchange Rate History,"
Working Papers w0210, Center for Economic and Financial Research (CEFIR).
Faberman, R. Jason, 2008. "Job Flows, Jobless Recoveries, and the Great Moderation," Federal Reserve Bank of Philadelphia Working Paper No. 08-11.
Pierce, Justin R., and Peter K. Schott, 2012. "The surprisingly swift decline of US manufacturing

Overall comments on referee #2. The key graf was "**In general, the approach and the findings of the paper are very interesting and compelling**. Yet I thought the paper was not always clear in terms of labeling, notation, and definitions." Naturally, I believe this was a fair assessment, and so in the current draft I try to be even more clear in terms of labeling, notation, and definitions, although given the lack of familiarity non-specialists have understanding the finer points of exchange rate indexing, some of these terms are always going to be a challenge for general economists. On the whole, this was a great referee report – thorough, and very helpful, with 90% of the comments on point, which is extremely high ratio by the normal standards of referee reports, and for someone who read the paper without being paid. I am thoroughly indebted to this referee.

QJE Reports (Submitted April, 2014)

Referee 1

Review of "Relative Prices, Hysteresis, and the Decline of American Manufacturing"

Summary

The paper examines the importance of real exchange rate movements in explaining the decline of US manufacturing. The author proposes that the strengthening of the dollar relative to foreign countries was responsible for the steep decline in manufacturing in the 1980s and 2000s, once a proper measure of the real exchange rate is employed. The paper suggests that a mechanism for these effects is through the mechanism of hysteresis, in which short-run strengthening of the dollar lead to job losses, and these jobs do not necessarily return even after the dollar weakens. As a separate case study, the author describes the Japanese experiment and suggests that it fits the paper's thesis: Japanese manufacturing employment shrank once the yen strengthened.

General Comments

I found this paper very enjoyable to read. It is a nice and thorough look at how real exchange rates have corresponded with changes in US manufacturing employment. I am not 100% convinced that one can treat these exchange rates as exogenous as the author does. For example, China's depreciation of the yuan was a coordinated policy move in conjunction with other changes, such as China's opening Special Economic Zones. The claim that the relative strengthening of the dollar *caused* the increase in exports to China trivializes the role of institutional changes which enabled US firms to exploit their access to cheap Chinese labor, which would have been cheap regardless of the relative strength of their currency. Or in the Japanese context, the strengthening of the yen was probably in part related to the economy reaching a certain maturity, and reaching its maximum net export potential. But by and large, I am a fan of the paper. It considers an interesting question in a thorough way. One last general point – the paper does feel a bit like 2 papers smushed together, one about the role of real exchange rates and another about hysteresis. It doesn't feel like it's too disconnected, but it's not clear the hysteresis point is really part of the same paper by necessity.

Response: Thank you for this report. The point that the Chinese government had an entire range of policies designed to promote industry, aside from an undervalued exchange rate, is an important one. To some extent, I think it is not possible to control for all of these. But I do think it is possible to argue that any general factors which raised productivity in China – such as subsidies to manufacturing -- would also have necessarily lowered China's RULCs. Also, the SEZs were opened between 1980 and 1984, while the US WARULC index appreciated between 1996 and 2003, and it was in this period (1997-2004, given the lags with which RERs operate) when most of the jobs were lost. Additionally, if this were all about China's trade policies, then the question is why didn't Canada experience a collapse at the same time as the US? China's policies would have had to have been US-centric, and also would have had to been implemented mostly in a 7-year period and then removed. In the Japanese case, the exchange rate movement was pretty sudden, whereas one would think that there would be a more gradual decline in export growth as a country meets its potential. It is also true that Japan never actually fully converged on the US standard of living/GDP per capita. However, these are fair points to raise and I will try to address this more in subsequent drafts.

Specific Comments

1. What is the number of observations in Table 1?

Response: Fixed!

2. I think all the figures need more notes.

Response: Good point. I have added in many more notes.

Conclusion

I think this paper merits serious consideration at a top journal like the QJE. While I think there are some holes in the argument that should perhaps be acknowledged more thoroughly, it tackles an important question and does so in an engaging way.

Referee 2

This paper examines the relationship between the decline in US manufacturing employment and movements in the US real exchange rate (RER) using an updated measure of the RER. It argues that exchange rate movements explain a substantial portion of job loss even after controlling for other explanations that have appeared in the literature, e.g., the growth of US-China trade.

1. The paper is longer and denser than it needs to be. I think it should start off with a clear description of the empirical strategy and analysis, and then use the model as needed to interpret the results. The large number of figures also makes the analysis a bit hard to follow. It might be better in terms of sharpening the focus on the main results of the paper to push some of them to the appendix while keeping their message in the main text.

Response: Good suggestion. I moved the model and several of the Figures to the appendix, and shortened the paper.

2. Appreciation of the dollar allows local manufacturers to procure imported intermediate inputs more cheaply, potentially increasing manufacturing activity. Is this channel relevant? Can it be accounted for in the empirical analysis?

Response: Note that I did control for intermediate inputs in Table I (formerly Table II). While I didn't get a generally significant result, I did get a significant negative results when I used a quantile regression. This suggests that cheap intermediates were used as substitutes for domestic production.

3. It would be interesting to see more about the updated components of the RER, especially by country.

Response: I did a bit of this – breaking out China from the rest (now in Table II), and also imports from exports. It's hard to do this generally for several reasons. First, note that the two periods of dollar appreciations, the dollar itself was generally moving around much more than most of the currencies of US trading partners, and so generally all of the bilateral RER indices are very highly correlated. Second, imports are also, of course, also very highly correlated by sector. Thus, if you wanted to interact the bilateral USA-Germany RER index with German import penetration, and separately interact the bilateral USA-France RER index with French Import penetration (or average of import penetration and export share), these two variables would likely be very highly correlated over time and cross-sectionally.

4. “Openness” is defined using the average of the share of exports in shipments and the share of imports in domestic consumption, so it would be interesting to split these out in the regression analysis. (In Figure IX, it would also be interesting to break down the components of openness and ULCs by country) However, these measures may be endogenous. Is it possible to U.S. trade policies as instruments, e.g., the variable about China's change in MFN status that also appears in the regression or the tariff changes of the Uruguay round that were implemented over this period? One would then ask whether movements in the RER amplify exogenous changes in U.S. trade flows. Also, lots of other trade policies were changing around the world during this period.

Response: Good suggestion. I have now moved this from the online appendix to the text of the paper. Note that I do try controlling for tariff changes, but the changes in US tariffs in this period were relatively small. It would be nice to have a good measure of manufacturing tariffs faced by the US in this period. I did try controlling for Chinese tariffs, and didn’t find any significant impacts. This wasn’t surprising, as only a very small share of US manufacturing goods are shipped to China.

E.g., in return for permission to enter the WTO, China implemented reforms that included lowering tariff rates, reducing subsidies and promising not to discriminate against foreign investment. Is it possible to control for policy changes like this or account for their effect on real exchange rates?

Response: Another good suggestion. I did try controlling for lower Chinese tariffs (as mentioned above). Reducing subsidies would presumably impact China’s ULCs. The WTO PNTR-NTR Gap variable (from Pierce and Schott) should presumably control for some of issues related to China’s WTO ascension.

5. Minor point: the paper also uses the term "tradability"; is that a synonym for "openness"?
Response: Yes it was. I have now just used openness for consistency.

6. Are the regressions behind Figure IV and V weighted? If not, shouldn't they be?
Response: Yes, they are weighted.

7. Regarding the discussion on page 21, it is possible that investors anticipating China's entry into the WTO would begin acting before it actually occurred.

Response: Sure, this is possible. But Pierce and Schott (2015) seem not to find much of an anticipatory effect. When they use annual dummies interacted with their China PNTR-NTR Gap variable, it seems to become significant immediately after China’s WTO ascension.

Furthermore, how related is the "opening" of China which occurred over this period to the US RER movements, given the way they are constructed?

Response: Of course, very related. Using my indexing method, growth in trade with a country with low prices/RULCs will, by definition, lead to an appreciation of the index. And I would argue would also lead to a decline in competitiveness. I also did a GDP weighted index, which would not at all be related to China’s opening but would be related to China’s fast GDP growth, and I get the same results.

8. What explains the asymmetry between appreciations and depreciations? Aren't US depreciations affected by other countries' appreciations?

Response: This indeed has been a big puzzle in the literature on US RER movements. I resolve it by noting that it is the level of the RER that matters, rather than the changes. If you appreciate 20%, but are still undervalued by 40%, then you will continue to accumulate tradable-sector jobs. If you appreciate by 20%, leaving you 40% overvalued, you will be losing tradable sector jobs. In fact, in almost any sort of economic model you write down, you'll get the implication that relative prices matter and not merely the change in relative prices. The short answer to the second question is that yes, of course, if the US depreciates another country appreciates.

9. Misc: The second y-axis is not labelled in Figure II.
Response: Indeed, it was not. Nice catch.

Overall Comments: This referee provided many very thoughtful comments that I used to improve the paper immeasurably. I am very much indebted.

Referee 3

Summary
This paper focuses on the role of real exchange rate fluctuation on the changes in the structure of manufacturing for the US. Using a differences in differences strategy, the paper argues that the appreciation of the real exchange rate may account for about 2/3 of the decline in manufacturing employment in the early 2000s.
The paper also argues that the effects of changes in the real exchange rate will generally take time to be felt on the US because trade responds more in the long-run than in the short-run. The paper develops a model of heterogeneous firms with sunk export costs to make this point.

Response: Note: the argument of hysteresis is closely related but not exactly the same as saying that changes in the RER will be felt more in the long-run than in the short. Of course, if you become permanently overvalued, then of course the impact will accumulate over time. But what the finding in the paper says is that if you are overvalued for just a year or two, your tradable-sector employment will shrink during this time period and then will stay shrunk even after your RER returns to fundamentals.

Understanding how trade affects the distribution of activity across sectors of the economy is an important issue and area of active research. While the topic is of general interest, I do have some concerns that should be addressed and discussed in the paper.
1. In some sense, the paper is just telling us the elasticity of substitution between imports and domestic goods has the right sign. There are many papers that already tell us that and also argue that the short-run elasticity is less than the long-run elasticity. The differential employment growth across sectors is just a corollary of these well-known empirical findings.

Response: Well, this paper also tells us that (1) historical prices matter – hysteresis is real, and (2) relative price movements and trade can explain most of the collapse in manufacturing employment in the

early 2000s. It is worth mentioning that RER movements have, to date, not been one of the factors cited in the literature so far about the decline in US manufacturing, so I think it is debatable whether this is that obvious.

2. I had a hard time reconciling the findings here with those in Ruhl, Steinberg, and Kehoe (2012) which basically find a relatively small effect of net exports and real exchange rate fluctuations for the shift out of manufacturing. The paper need to relate to that work.

Response: Actually, I would liken their calibration approach to the accounting approach in Table VIII. The findings are actually the same, just spun differently. What Kehoe *et al*. say is that, had manufacturing productivity suddenly stopped after 2000, that, in a model where productivity and employment in manufacturing have a one-for-one negative relation, that there would have been no decline in manufacturing employment. My accounting approach suggests the same. Note that their method also implies that had productivity suddenly stopped after 1980, manufacturing employment would have thereafter grown substantially. However, they only start their analysis in the 1990s. Another puzzle is why, if a huge boom in US manufacturing productivity caused the decline in employment, why there also emerged such a large trade deficit. Generally, productivity growth should cause a trade surplus. Another problem is that, as mentioned in the paper, while measured productivity growth after 2000 was the same as it was before 2000, Houseman et al. (2010) have highlighted an index numbers problem with the official numbers after 1997, and suggest up to one-fourth to one-half of this productivity growth might not have actually happened. Lastly, most of the productivity growth in this period happened in one industry (computers) – productivity in the median manufacturing sector actually *declined* in the 2000s. Even stranger, total sales in the computer sectors actually declined (only, employment declined by even more), and this is the sector in which the index numbers problem looms the largest. None of these facts are considered in Kehoe *et al*., and none are consistent with the rosy picture they paint.

3. I do not see the need to motivate the dynamic relationship between the real exchange rate and net exports with the model presented. It is well accepted that sunk export costs can alter the lag-lead relationship between a change in the real exchange rate (Baldwin & Krugman, Roberts & Tybout, Das, Roberts& Tybout). The key question is whether this is a strong enough channel relative to other types of frictions. The GE work by Alessandria & Choi (QJE, 2007) and Alessandria, Pratap, & Yue (2011) explore these ideas theoretically and empirically and find some ability of this mechanism to account for the data. I think you would just be better off appealing to this earlier work than deriving your model.

Response: I agree with this. I now have moved the model to the appendix.

4. The paper should really relate to the open economy structural change literature (see Uy, Yi, Zhang or Bett and Giri or many others). A key element of that work says that comparative advantage can accelerate structural change in advanced countries. This work tends to focus on balanced trade, but it clearly points out that the open economy may be important for the pace of structural change.

Response: These are interesting papers, but I do not seem them as being that closely related to the task at hand.

5. I am a little worried that the paper does not seem to properly take trends in the data into account. For instance, Figure 1 shows fairly clearly to me that there is a fairly constant decline in share of manufacturing employment in the population. There are certainly wiggles around this trend as the pace of decline in manufacturing seems slow in booms and accelerate in busts. These fluctuations are quite consistent with manufacturing intensively producing cyclical goods like capital and durables.

Response: I explicitly control for trends in all regressions in the paper. I also plot pre-treatment, and post-treatment trends, so it can be seen that trends are not driving my results. In addition, I even try controlling for trends-in-the-trends – the second derivative of sectoral employment growth, and the results are little changed. I also try lagging the dependent variable, and the results hold. Also, while it is admittedly hard to see, in Figure 1, after the 2000s, manufacturing employment as a share of the population is below trend. It's also the case that if you extrapolate a linear trend another 30 years or so, it will become negative. Thus, even if it were at a linear trend, you would expect it to flatten like agriculture has done.

a. Incidentally, it would be much more useful to scale manufacturing employment by overall employment.

Response: I disagree. See the footnotes on page1.

6. In some sense the empirical results about more open sectors growing slower than less open sector following an appreciation have to hold given the aggregate data. What I mean is that we know that net exports falls with a lag following an appreciation. The only way the results could go in a different way is trade liberalization went in an opposite way.

Response: I would agree that the results are intuitive. And yet, I would also argue that the implicit model even most trade economists have in their head about the way the world works does not incorporate a role for hysteresis. Also, the finding that it is the level of relative prices and not the log change that matters apparently is not very intuitive given your next comment.

7. The paper seems to find an asymmetry in the effect of exchange rate movements on international re-allocation with appreciations have a negative effect on more open industries and a depreciation having no beneficial effect. This asymmetry is puzzling. I don't see how the mechanism can only work in one direction.

Response: This is explained in the text of the paper, on page 6: "One puzzle is that RER appreciations have been noted to lead to declines in employment, while depreciations do not seem to help. The solution to the puzzle is that what matters for employment is the level of US relative unit labor costs, rather than the change. The reason is conceptually easy to understand -- if unit labor costs were the same in the US and China, there would be no economic reason to move production, particularly as this could entail substantial fixed costs. On the other hand, when US unit labor costs are 50% higher than in trading partners, there is clearly an economic incentive to shift production, while firms already located abroad would have a competitive advantage. This finding should not be suprising in light of the central tenet of economics, that prices matter."

Overall comments: This referee made many useful suggestions which I have used to improve the paper.

Referee 4

This paper considers the contribution of international factors to manufacturing employment in the United States. The author states that the focus of this paper is to determine "…how much of the collapse in manufacturing in the early 2000s can be explained by relative prices." (p. 2) The analysis considers a set of episodes (the US in the 1980s and the early 2000s, Canada in the middle of the 2000 – 2010 period, and Japan in the 1980s). The author makes the case that the appropriate indicator for conducting an analysis like this one, for the United States, is the Weighted Average Relative Unit Labor Cost Index (WARULC) rather than the more commonly used broad trade-weighted real exchange rate index calculated by the Federal Reserve. The main result from this paper is that an exchange rate appreciation adversely affects manufacturing employment, but there is less evidence that a depreciation raises manufacturing employment.

This result is interesting, but results like this are already in the literature. Therefore, a more appropriate outlet for this paper would be a field journal, rather than a high-level general interest journal like the QJE. But before the author resubmits this paper to a field journal, I suggest an extensive re-write.

Response: If there is a well-identified paper on exchange rate hysteresis using disaggregated data, as opposed to using trends in Macro evidence or a calibrated model, I would like to see it. I would also like to see a paper explicitly linking movements in relative prices to the decline in manufacturing in the early 2000s. Once again, the collapse in US manufacturing, according to people like Ben Bernanke, was one of the seminal events in the past 50 years of US economic history. To my knowledge, no one has written about the link between relative prices movements and the collapse in US manufacturing, even though there have been many papers linking trade and the rise of China to the collapse (published in top journals as well). The estimates I get for the direct impact of RER movements and trade on the manufacturing sector in this period are more than five times larger than the Acemoglu *et al*. (2015) estimates. My estimates could be wrong or right, but it seems strange to argue that the sudden loss of 3 million manufacturing jobs isn't important. I had worried that my conclusions were, if anything, too provocative.

Another thing to note is that even if my paper were just about the short-run impact of RER movements, I believe I make plenty of improvements over what came before. Klein *et al*. (2003) is the most closely related literature, and while it is an excellent paper, it came in the era before single-way clustering, much less multi-way clustering with quantile regressions. I use much more data, more than twice as many control variables, fix an index numbers problem in the RER indices they use, utilize repeated out-of-sample testing, and at least make every attempt to think hard about identification. I also solve the puzzle they highlight about the asymmetry of RER movements and suggest a new functional form for how relative prices influence the economy. But the big contribution with first-order policy importance is the finding of hysteresis.

There are a number of conceptual issues that need to be addressed in this paper. One central issue is why the author considers only subsets of the sample period.

Response: I consider the full period from 1972 to 2009 in my regression results. The subsets are merely for the illustrations.

The argument is made that this is important to ensure that the real exchange rate is not being driven by activity in the manufacturing sector, but my sense is that at high frequencies this is notreally a concern, especially given the limited size of the manufacturing sector. At a minimum, if the author wants to continue to use the subsample approach, elasticity estimates using these observations should be compared those estimated with the full sample. Also, it is not clear to me why the falsification exercise (p. 27) should make me more confident in the results of the paper.

Response: The point of the falsification exercise is to show that the estimation method is not prone to generating spurious results.

Finally, I did not understand how the input-output linkages used to account for job losses (p. 29) are not already accounted for with regression estimates of elasticities with respect to the real exchange rate.

Response: Imagine a factory faces a flood of cheap imports, and closes down as a result. Then that factory will stop buying intermediate goods from local producers, and so this will be recorded as a decline in demand for those sectors. This may be a reason why observed demand growth appeared sluggish in the 2000s.

The paper also needs to be rewritten in a way that makes it more "linear." There are many discussions in the paper that can be viewed as "asides." The paper requires extensive editing, eliminating much of what appears in the current version. The theory section can be cut – it is not needed to motivate the empirical analysis and, indeed, there is no explicit exchange rate in the model (as mentioned on p. 10) so its usefulness in this paper is limited. Many of the statistical presentations do not bear directly on the question at hand. For example, much of Section 2 up to Section 2.2, can be cut – why show univariate relationships when you will present multivariate regressions that are much more informative? You don't need Figures VI, VII and X to motivate the empirical analysis (as stated on p. 20) since the specification is standard. There should be a more extensive discussion, however, on the comparison of the WARULC index and other indices, since this seems to be a central issue in this paper; what were the sources of divergence across indices? There is some discussion of this on pp. 25 – 26, but this discussion should be expanded and moved into the Data Section.

Response: These are solid points. I eliminated or moved some of the Figures, and moved the model to the appendix, and moved the introduction of WARULC to the data section.

Overall Response: This referee made some sharp suggestions to which I am very grateful.