A Report on "Trade with Nominal Rigidities: Understanding the Unemployment and Welfare Effects of the China Shock" by Rodríguez-Clare et al. (2026)

Reviewer 2

February 22, 2026

v1

isitcredible.com

Disclaimer

This report was generated by large language models, overseen by a human editor. It represents the honest opinion of The Catalogue of Errors Ltd, but its accuracy should be verified by a qualified expert. Comments can be made here. Any errors in the report will be corrected in future revisions.

> I am wiser than this person; for it is likely that neither of us knows anything fine and good, but he thinks he knows something when he does not know it, whereas I, just as I do not know, do not think I know, either. I seem, then, to be wiser than him in this small way, at least: that what I do not know, I do not think I know, either.
>
> Plato, *The Apology of Socrates*, 21d

> To err is human. All human knowledge is fallible and therefore uncertain. It follows that we must distinguish sharply between truth and certainty. That to err is human means not only that we must constantly struggle against error, but also that, even when we have taken the greatest care, we cannot be completely certain that we have not made a mistake.
>
> Karl Popper, 'Knowledge and the Shaping of Reality'

Overview

Citation: Rodríguez-Clare, A., Ulate, M., and Vásquez, J. P. (2026). Trade with Nominal Rigidities: Understanding the Unemployment and Welfare Effects of the China Shock. *Journal of Political Economy*, Vol. 134, No. 2, pp. 626–664.

URL: https://doi.org/10.1086/738344

Abstract Summary: This paper develops a dynamic quantitative trade and migration model incorporating downward nominal wage rigidities to analyze the unemployment and welfare effects of the China shock. It finds that while the China shock increases welfare in most US states, nominal rigidities significantly reduce these gains and lead to welfare losses in some states.

Key Methodology: Dynamic quantitative trade and migration model with downward nominal wage rigidities (DNWR), calibrated using dynamic hat algebra to match empirical evidence from Autor, Dorn, and Hanson (2013).

Research Question: How do downward nominal wage rigidities affect the unemployment and welfare outcomes of trade shocks, specifically the China shock?

Summary

Is It Credible?

Rodríguez-Clare et al. develop a dynamic quantitative trade and migration model to evaluate how downward nominal wage rigidities (DNWR) shaped the US labor market's response to the "China shock." The authors claim that while the China shock improved the terms of trade for the United States, the inability of nominal wages to adjust downward quickly led to temporary but significant unemployment. Their headline quantitative finding is that DNWR reduced the overall US welfare gains from the China shock by roughly two-thirds, dropping the aggregate gain from 31 basis points to 12 basis points, and causing 18 states to experience welfare losses that would have otherwise seen gains (pp. 626, 629). The authors also argue that their framework successfully replicates the spatial heterogeneity and persistent employment declines documented in foundational empirical work.

The credibility of the precise two-thirds welfare reduction is limited by the extreme sensitivity of the model to the assumed duration of the trade shock. The baseline model assumes the China shock lasted only from 2001 to 2007. However, the authors acknowledge that Chinese import penetration continued to grow until around 2010. When they run an alternative "longer shock" specification lasting until 2011, the aggregate welfare gains for the US "roughly disappear," falling to just 0.011 percent (p. 630; Table 1, p. 649). The authors truncate the baseline shock at 2007 to avoid mechanical complications related to the Great Recession, but this exposes a foundational fragility in the analysis: the model assumes the China shock was the only economic disturbance affecting the US economy during this period (p. 647, fn. 23). While the authors argue in Section VIII.B that concurrent events like the housing bubble might have actually masked the full impact of the trade shock, structurally omitting these complex macroeconomic dynamics risks misattributing labor market

outcomes entirely to the specified trade transmission mechanism (p. 659).

Furthermore, the model's internal mechanics present several inconsistencies that complicate its interpretation. The authors motivate the article with new empirical evidence showing that cross-state heterogeneity in DNWR stringency is a significant predictor of differential unemployment responses (p. 635). Yet, the quantitative model assumes a single DNWR parameter (δ) that applies to all manufacturing sectors in all US states (p. 647). Although the *incidence* of rigidity varies by region based on shock exposure, and the authors test sensitivity to this assumption, the structural choice to use a uniform parameter sits somewhat uneasily alongside the empirical motivation (p. 657). Additionally, the model relies on a stylized nominal anchor—assuming world nominal GDP in US dollars grows at a constant rate—to force the wage rigidities to bind (p. 628, fn. 4; p. 642). The authors justify this as capturing global aggregate demand constraints, but without an active, realistic monetary policy response targeting domestic inflation and unemployment, the model arguably maximizes the disruptive impact of DNWR.

The calibration strategy also relies on strong assumptions about the nature of the unemployment generated by the China shock. The model is calibrated to match the employment and population impacts estimated by Autor et al. However, the model treats all generated unemployment as purely cyclical (p. 643, fn. 15). While the authors defend DNWR as a mainstream tool for explaining involuntary unemployment during transitions, mechanically forcing nominal frictions to explain empirical moments that likely contain deep structural components—such as skills mismatch or the unwinding of local agglomeration economies—may overstate the explanatory power of DNWR (p. 632). The authors acknowledge that their framework omits mechanisms like human capital depreciation and hysteresis (p. 661).

Finally, the assumption of homogeneous workers within any given sector-region limits the model's ability to capture the true distributional pain of the China shock. Because all workers in a sector face the same wage and employment risks, the model

cannot address the concentrated losses experienced by lower-wage workers, a dynamic well-documented in the broader literature. While the model struggles to match untargeted wage effects—overpredicting the fall in nonmanufacturing wages by more than 50 percent compared to empirical estimates—the authors note this is at least qualitatively consistent with the data (Table 1, p. 649; p. 650). Ultimately, the article is a valuable theoretical exercise. It successfully demonstrates the qualitative mechanism by which nominal rigidities can transform a positive terms-of-trade shock into localized unemployment, even if the exact quantitative welfare estimates are too dependent on rigid structural assumptions to be taken as definitive.

The Bottom Line

The article provides a compelling theoretical framework demonstrating how downward nominal wage rigidities can transform a positive terms-of-trade shock into localized unemployment. However, its headline quantitative claim—that these rigidities erased two-thirds of the US welfare gains from the China shock—is highly sensitive to structural assumptions. Specifically, the welfare estimates are fragile to the assumed duration of the shock, the exclusion of concurrent macroeconomic events like the Great Recession, and the use of a stylized nominal anchor. While the qualitative mechanism is plausible and represents a valuable contribution to the trade literature, the precise welfare and employment estimates should be viewed with considerable caution.

Potential Issues

Sensitivity of welfare results to the shock's duration: The article's headline quantitative results, including the central claim that Downward Nominal Wage Rigidity (DNWR) reduces welfare gains by two-thirds, are derived from a baseline model where the China shock is assumed to last from 2001 to 2007. The article itself presents an alternative "longer shock" specification (lasting until 2011) that it notes is more consistent with recent empirical work and provides a much better fit to the dynamic persistence of the employment effects observed in the data (p. 656). The welfare implications of this better-fitting model are starkly different: the aggregate welfare gains "roughly disappear," falling from 0.126% in the baseline to just 0.011% (p. 630; Table 1, col. 3, p. 649). The authors justify the choice of the shorter baseline to "avoid complications related to the Great Recession," as falling trade during that period would be mechanically interpreted by the model as a negative productivity shock in China (p. 657). While the authors are transparent about this choice and its consequences, the extreme fragility of the main welfare conclusion to the shock's end date is a critical structural feature of the analysis.

Foundational assumption of a single economic shock: The model's identification and calibration of its key parameters—downward wage rigidity (δ), inter-sectoral mobility (ν), and inter-regional mobility (κ)—depend on the assumption that the China shock was the only economic disturbance affecting the US economy between 2000 and 2007. The article acknowledges this in a footnote, stating that "Identification relies on the assumption that the China shock is the only shock affecting the model economy" (p. 647, fn. 23). While the calibration strategy uses an instrumental variable approach to target the causal estimates from Autor et al., thereby isolating the shock from raw correlations with concurrent events, the general equilibrium counterfactual still assumes no other shocks are interacting with the China shock. The authors do address this limitation, arguing that concurrent shocks like

the housing bubble likely kept unemployment artificially low, meaning their results might be conservative (p. 659). However, by structurally omitting these concurrent shocks from the model's dynamics, the analysis may incorrectly attribute all observed labor market changes to the specified transmission mechanism.

Dependence on a stylized nominal anchor: The model's ability to generate unemployment relies on a specific nominal anchor: the assumption that world nominal GDP in US dollars grows at a constant rate, which is then normalized to zero for the calibration (p. 628, fn. 4; p. 642). This assumption is a critical driver of the results, as a negative labor demand shock in a fixed nominal income world necessitates a fall in nominal wages, which then triggers unemployment if DNWR is present. The authors defend this choice as capturing a "given level of world aggregate demand in the context of a global savings glut" and argue that alternative nominal assumptions yield qualitatively similar results (pp. 642, 659). They also explore the "sacrifice ratio" in Section VIII.C, finding a non-linear relationship where small amounts of inflation can reduce unemployment, but further reductions become costly (p. 660). Nevertheless, the baseline model lacks a realistic, active monetary policy response, which remains a significant abstraction.

Omission of alternative mechanisms for labor market persistence: The article frames DNWR as the central mechanism explaining the persistent employment effects of the China shock. However, the model does not incorporate or test against other well-established mechanisms that could also generate such persistence, such as slow capital adjustment, the depreciation of sector-specific human capital, or the unwinding of agglomeration economies. The article acknowledges this limitation in its concluding paragraph, noting that the model "does not incorporate mechanisms such as human capital depreciation, hysteresis, or agglomeration" (p. 661). By focusing exclusively on DNWR and a specific labor mobility structure, the model may attribute the entirety of the observed persistence to these channels, potentially resulting in an upward bias in their estimated importance.

Assumption of homogeneous workers: The model assumes that all workers within a given sector-region are homogeneous, earning the same wage and facing the same employment risks. This is a significant simplification that prevents the analysis from addressing the differential impacts of the China shock on workers of different skill or wage levels, a key distributional channel highlighted in the empirical literature. The authors explicitly acknowledge this limitation in the conclusion, stating that their approach is "inconsistent with evidence... that lower-wage workers... experience worse earnings trajectories" (p. 661). Consequently, the article's headline finding of small positive average welfare gains may mask a more complex reality of concentrated losses for the most vulnerable workers, a dimension the model is structurally unable to capture (p. 652).

Internal inconsistency in the treatment of wage rigidity: The article presents an internal tension between its motivating empirical evidence and its quantitative model. In Section II.C, the authors provide new empirical evidence suggesting that cross-state heterogeneity in the degree of DNWR is a significant predictor of differential unemployment responses to the China shock (p. 635). However, the quantitative model used for the main analysis assumes a single DNWR parameter (δ) that "applies to all manufacturing sectors in all US states" (p. 647). While the authors note that the *incidence* of rigidity varies by region and test the sensitivity of the uniform parameter assumption, the model's structure is technically inconsistent with the article's own empirical findings, which suggest that heterogeneity in the parameter itself is a key explanatory variable (p. 657).

Calibration to potentially contested empirical targets: The model's key parameters are calibrated to match three specific regression coefficients from Autor et al. The article acknowledges in a footnote that the statistical validity of these foundational targets has been questioned in subsequent research, citing Borusyak et al. who "have cast doubt on the statistical significance of some of the results" (p. 647, fn. 22). The authors justify their choice based on the influence of the original article, stating they

aim to provide a structural interpretation of these "well known" estimates. However, they do not provide a sensitivity analysis to show how the results would change if calibrated to smaller or statistically insignificant targets. This approach risks building a complex quantitative structure to explain what might be statistical noise.

Discrepancy in untargeted wage effects: While the model is calibrated to match employment and population moments, its predictions for untargeted wage changes differ from the empirical estimates in Autor et al. The baseline model predicts a nearly zero change in manufacturing wages (0.023) compared to the empirical estimate of 0.150, and it over-predicts the fall in nonmanufacturing wages by more than 50% (-1.177 vs. -0.761) (Table 1, p. 649). The authors argue this is "qualitatively consistent" because nonmanufacturing wages fall more than manufacturing wages in both the model and the data (p. 650). However, the quantitative discrepancy for nonmanufacturing wages is notable. Furthermore, the empirical estimate for the manufacturing wage change is not statistically significant in the original study, which makes the model's near-zero prediction less of a clear mismatch with the data than the raw point estimates suggest.

Mismatch between geographic units of analysis: The model is calibrated and analyzed at the US state level, whereas the empirical targets from Autor et al. are estimated at the more granular level of local labor markets, or commuting zones (CZs). The authors acknowledge this discrepancy and justify the state-level analysis based on the availability of bilateral trade data. Crucially, they also report that running the original regressions at the state level "yields similar response-to-exposure coefficients," providing an empirical validation for their choice of geographic unit (p. 649, fn. 26). Nonetheless, this aggregation may mask significant intra-state heterogeneity and represents a potential misalignment between the economic mechanisms, which are understood to operate at the local labor market level, and the model's state-level structure, which implicitly assumes perfect labor mobility within states.

Asymmetric assumptions on international labor mobility: The model assumes per-

fect, costless sectoral labor mobility for all countries outside the United States, stating that "there are no costs of moving across sectors within a region" (p. 645). This stands in stark contrast to the carefully calibrated, friction-filled mobility structure for US states. The authors note this simplification is made because constructing mobility matrices for 36 other countries "would be extremely cumbersome" (Supplementary Materials, p. 36). Crucially, however, the model assumes flexible exchange rates for these countries, which allows them to maintain full employment regardless of mobility frictions (p. 641). While this mitigates the direct impact of the mobility assumption, the asymmetry in structural detail between the US and the rest of the world remains a notable feature of the design.

Conceptual mismatch in the definition of unemployment: The model is calibrated to match the empirical effect of the China shock on the unemployment-to-population ratio, as measured by Autor et al. The authors state that "the concept of unemployment in our model is that of cyclical unemployment, i.e., the unemployment in excess of the natural rate" (p. 643, fn. 15). This creates a potential conceptual mismatch. The China shock likely had structural effects that could have increased the natural rate of unemployment (e.g., through skills mismatch), not just the cyclical component. The model targets the estimated change in unemployment (a regression coefficient of 0.22 percentage points), not the total unemployment rate (p. 647). However, by forcing a model of purely cyclical unemployment to match an empirical moment that may include both cyclical and structural components, the calibration may mechanically attribute structural job losses to nominal frictions, potentially overstating the role of DNWR.

Future Research

Heterogeneous wage rigidities: Future work could incorporate spatially and sectorally varying downward nominal wage rigidity parameters into the quantitative model, aligning the structural framework with the empirical evidence that regional differences in wage stickiness drive differential unemployment responses.

Incorporating structural unemployment and hysteresis: Researchers could extend the model to include search and matching frictions, human capital depreciation, or agglomeration effects. This would allow the framework to distinguish between the purely cyclical unemployment caused by nominal rigidities and the structural, long-term non-employment that characterized the hardest-hit local labor markets.

Endogenous monetary policy: Future models could replace the stylized constant global nominal GDP anchor with an active central bank that follows a Taylor rule. This would allow for a more realistic assessment of how domestic monetary policy might offset or exacerbate the localized unemployment effects of trade shocks in the presence of nominal rigidities.

Worker heterogeneity and distributional impacts: Building models that allow for heterogeneous skill levels and wages within sector-regions would enable researchers to capture the disproportionate earnings and employment losses suffered by low-wage workers, providing a more accurate accounting of the shock's true welfare costs.

isitcredible.com