

Federal Reserve Bank of Minneapolis
Research Department Staff Report 489

Revised: October 2016
(First version: August 2013)

Global Imbalances and Structural Change in the United States*

Timothy J. Kehoe
University of Minnesota,
Federal Reserve Bank of Minneapolis,
and National Bureau of Economic Research

Kim J. Ruhl
Stern School of Business, New York University

Joseph B. Steinberg
University of Toronto

ABSTRACT

Since the early 1990s, as the United States borrowed heavily from the rest of the world, employment in the U.S. goods-producing sector has fallen. We construct a dynamic general equilibrium model with several mechanisms that could generate declining goods-sector employment: foreign borrowing, nonhomothetic preferences, and differential productivity growth across sectors. We find that only 15.1 percent of the decline in goods-sector employment from 1992 to 2012 stems from U.S. trade deficits; most of the decline is due to differential productivity growth. As the United States repays its debt, its trade balance will reverse, but goods-sector employment will continue to fall.

Keywords: Global imbalances; Real exchange rate; Structural change

JEL classification: E13, F34, O41

*We thank David Backus, David Strauss, Frank Warnock, Kei-Mu Yi, Vivian Yue, Jing Zhang, and four anonymous referees for helpful discussions and suggestions. Kehoe and Ruhl gratefully acknowledge the support of the National Science Foundation under grant SES-0962993. We are grateful to Jack Rossbach for extraordinary research assistance. The data and computer code used here can be found in the supplementary materials available at www.econ.umn.edu/~tkehoe. The views expressed herein are those of the authors and not necessarily those of the Federal Reserve Bank of Minneapolis or the Federal Reserve System.

1. Introduction

Between 1992 and 2012, the U.S. net international investment position deteriorated by \$4 trillion as households and the government in the United States borrowed from the rest of the world. During this same period, the goods-sector share of total employment fell sharply. Economic theory predicts that trade deficits drive employment out of domestic goods-producing sectors, and a commonly held view in policy circles, expressed, for example, in Scott (2015), is that U.S. trade deficits generated by this borrowing have played an important role in the decline in U.S. goods-sector employment and that rebalancing U.S. trade will reverse a large part of this trend. In this paper, we provide the first model-based estimates of the impact of trade deficits on U.S. goods-sector employment.

We use a dynamic general equilibrium model of the United States and the rest of the world to address the questions: To what extent are trade deficits responsible for the loss of U.S. goods-sector employment? Will employment return to goods-producing sectors when U.S. borrowing ends and trade deficits become trade surpluses?

Our framework combines an open-economy model of intertemporal trade with a model of *structural change* — the secular shift of employment from goods-producing industries to services-producing industries. Our model features two main ingredients from the structural change literature: income effects from nonhomothetic preferences, and faster labor productivity growth in the goods-producing sector compared with the services sector combined with a low elasticity of substitution between goods and services. Typically, structural change has been studied in closed-economy settings; we extend the analysis to an open economy with unbalanced trade.

The key open economy feature in our model is the *saving glut* — increased demand for saving by the rest of the world — which we calibrate to match the U.S. trade balance. We use our model to quantitatively assess the contributions of traditional structural change forces and the saving glut to the decline in U.S. goods-sector employment. Contrary to assertions by analysts such as Scott (2015), who writes that “trade [deficits], not productivity, is the culprit [behind goods-sector job loss],” we find that the saving glut only explains between 11 and 20 percent of the decline in U.S. goods-sector employment; our preferred estimate is 15.1 percent.

We calibrate our model’s production and demand structure so that the model replicates the 1992 U.S. national accounts and a world input-output matrix we construct from the World Input-Output Database (Timmer et al., 2015). We feed into the model two exogenous driving forces:

labor productivity growth rates that differ across sectors and demand for saving in the rest of the world. The equilibrium outcome of this model — our *benchmark saving-glut equilibrium* — accounts for 99.2 percent of the observed decline in the goods-sector employment share from 1992 to 2012 and accounts for several other key facts about the U.S. economy during this period.

We then compare this equilibrium with one in which foreign demand for saving is constant — our *no-saving-glut counterfactual*. The counterfactual accounts for 84.1 percent of the decline in the goods-sector employment share. The difference between these two figures, 15.1 percent, is the contribution of the saving glut to this decline. This implies that eliminating the trade deficit will not generate a significant increase in goods-sector employment.

In figure 1, we plot the share of goods-producing sectors — agriculture, mining, and manufacturing — in total employment, which has fallen as the trade deficit has grown. The economic theory that links these data depend on imported goods being substitutable for domestically produced goods. As the United States trades bonds for foreign goods, labor shifts away from domestically produced goods and toward services and construction, which are less substitutable for foreign goods. This mechanism implies that, when the debt has to be repaid, labor will flow back into the U.S. goods sector to produce the extra goods needed to repay the debt. The quantitative importance of this channel in driving the decline in U.S. goods-sector employment over the past two decades is the focus of our study.

There is ample historical evidence connecting sectoral reallocation and trade balance dynamics. After Spain entered the European Community in 1986, its trade balance deteriorated substantially, while the goods-sector share of Spanish GDP fell from 40 percent to less than 30 percent (Fernandez de Cordoba and Kehoe, 1999). The Baltic countries had similar experiences after liberalizing capital markets in the 1990s (Bems and Jönsson, 2005). In emerging economies like Mexico, “sudden stops” of the world’s willingness to lend are typically accompanied by sharp reallocations away from services production and into goods (Kehoe and Ruhl, 2009). Although the saving glut is not as drastic as capital account liberalization or reversals due to sudden stops, it is plausible that 20 years of increased capital inflows could have a significant effect on goods-sector employment.

In constructing our model, we need to specify the driving force behind U.S. borrowing. A common explanation is that foreign demand for saving increased, making foreigners more willing to trade their goods for U.S. bonds. Bernanke (2005) coined the term *global saving glut* to refer to

this mechanism, and we adopt Bernanke's global-saving-glut hypothesis. Several explanations have been proposed for the increased demand for saving in the rest of the world, such as a lack of financial development in the rest of the world (Caballero et al., 2008; Mendoza et al., 2009) and differences in business cycle or structural growth properties (Backus et al., 2006; Perri and Fogli, 2010; Jin, 2012). We do not take a stand on which of these explanations are correct. Instead, we take the saving glut as given and calibrate a process for household preferences in the rest of the world over current versus future consumption so that our model matches exactly the path of the U.S. trade balance during 1992–2012.

The driving force behind structural change in the model is differential labor productivity growth. As evident in figure 2, labor productivity in the goods sector has grown at a faster pace than in other sectors over the past two decades. The structural change literature emphasizes asymmetric productivity growth as an important driver of the long-run reallocation of labor across sectors (Ngai and Pissarides, 2007). We take a similar approach in an open-economy model. Several other recent papers study structural change in open economies (Echevarria, 1995; Matsuyama, 1992, 2009; Sposi, 2012; Uy, Yi, and Zhang, 2013). With the exception of Sposi (2012), these studies use models of balanced trade, abstracting from international capital flows. Our model enriches the dynamics in Sposi (2012) by distinguishing between household and government debt, endogenizing the investment rate, and allowing countries to hold debt on the balanced growth path.

Buera and Kaboski (2009) identify three sources of structural change, all of which we include in our baseline model: differential sectoral productivity growth and a low elasticity of substitution between sectors (Ngai and Pissarides, 2007); nonhomothetic preferences (Kongsamut et al., 2001); and capital deepening driven by different capital shares across sectors (Acemoglu and Guerrieri, 2008). We find that differential capital shares are not significant sources of structural change in our model. Nonhomothetic preferences account for 5.6 percent of structural change, consistent with Buera and Kaboski (2009), who argue that, because subsistence requirements are most important at low income levels, income effects are unlikely to be important in explaining U.S. structural change in recent decades. Boppart (2014) and Herrendorf et al. (2013) find larger roles for income effects over longer time horizons.

We include several other features in our model that make it well-suited to address our questions. First, it includes three sectors: goods, services, and construction. Goods and services in

our model are both traded, in contrast to the usual convention in international macroeconomics which assumes only goods are tradable. The usual convention is at odds with the data: services are a large component of U.S. exports. Construction is the only purely nontradable sector and is used entirely to produce investment goods, which means that construction is more sensitive than the other sectors to the effects of capital flows and economic fluctuations in general. Second, our model includes a global input-output structure. Firms use intermediate inputs of both goods and services, produced both at home and abroad, and we distinguish between trade in final versus intermediate goods and services. This production and demand structure allows for substitution between domestic products and their foreign equivalents, and complementarity between products from different sectors. We also allow for more substitution between domestic and foreign goods than services, which allows us to match the volatilities in the goods trade balance and the services trade balance.

2. The U.S. economy, 1992–2012

We view the massive foreign borrowing and the differences in productivity growth across sectors as exogenous driving forces that we take as inputs into the model. We present three key facts about U.S. data that we view as tests for our model. Our model’s ability to capture these facts in response to the exogenous driving forces gives us confidence in using the model to perform counterfactual experiments and make predictions about the future of the U.S. economy.

Exogenous driving force 1: Foreign borrowing increased, then decreased

Figure 3 illustrates how much borrowing households and the U.S. government have done. The current account balance measures capital flow into the United States, but we see that the trade balance tracks the current account balance almost exactly. Since our model is not designed to accurately capture the difference between these two series, we use the trade deficit as our measure of U.S. borrowing. Figure 3 shows that between 1992 and 2006, the trade deficit grew steadily, reaching 5.8 percent of GDP, after which it began to shrink. In 2012, the trade deficit was 3.6 percent of GDP. We view the path of the trade deficit as what defines the saving glut in our model.

Capital flows are important to our analysis in two ways. First and foremost, the U.S. trade deficit, which is composed primarily of goods (figure 6), directly affects the need for labor to produce goods domestically. Second, accumulated debt eventually needs to be repaid, which could affect the employment needed to produce goods in the United States in the future. In our model, cumulative trade balances are the measure of this debt, which differs from the U.S. net foreign

asset position by any revaluation effects. These revaluation effects have, at times, played a significant role in the value of the U.S. net foreign asset position (Lane and Milesi-Ferretti, 2008; Gourinchas and Rey, 2007). Revaluation has been in favor of the United States overall; the U.S. net foreign asset position is smaller than cumulated trade balances in the data. This implies that future reallocation of labor back to goods-producing sectors caused by debt repayment may be smaller than our model predicts.

Exogenous driving force 2: Labor productivity grew fastest in the goods sector

During 1992–2012, labor productivity in the goods sector — real value added per worker — grew at an average of 4.4 percent per year, while it grew by only 1.3 percent per year in services and fell by 0.84 percent per year in construction. What is essential in our model is the differential between productivity growth in goods compared with services. As the data in figure 2 show, this differential has been close to constant since 1980, with average productivity growth of 4.3 percent per year in goods during 1980–2012, compared with 1.1 percent in services. Except for the productivity slowdown of the 1970s, the differential has persisted since 1960.

Fact 1: The real exchange rate appreciated, then depreciated

Figure 3 presents the first key fact that we ask our model to match. The figure shows that the U.S. real exchange rate was volatile and tracked the trade balance closely during 1992–2012. We construct our measure of the U.S. real exchange rate by taking a weighted average of bilateral real exchange rates with the United States’ 20 largest trading partners, with weights given by these countries’ shares of U.S. imports in 1992 (other weighting schemes yield similar results). The real exchange rate appreciated by 27.9 percent between 1992 and 2002, after which it depreciated by 22.1 percent between 2002 and 2012.

The intuition for why the real exchange rate and trade balance should move closely is straightforward: as foreign goods and services become cheaper, U.S. households buy more of them. Notice, however, that the timing is off: the maximum appreciation of the real exchange rate occurred in 2002, whereas the largest trade deficit occurred in 2006. Our baseline model is unable to replicate this pattern; we revisit it in our sensitivity analysis and concluding remarks.

Fact 2: The goods sector drove aggregate trade balance dynamics

Figure 6 presents our second key fact, plotting disaggregated trade balances for goods and services separately. The goods trade balance generates most of the fluctuations in the aggregate trade

balance, while the services trade balance fluctuates between 0.5 and 1.2 percent of GDP. The consistent U.S. services trade surplus motivates one of the key features of our modeling framework. Standard modeling conventions in international macroeconomics lump services together with construction into a nontradable sector, treating goods as the only sector that produces output that can be traded internationally. By contrast, we allow both goods and services to be traded in our model.

Fact 3: Employment in goods declined steadily; construction employment rose, then fell

Figure 7 presents our third fact: between 1992 and 2012, the fraction of total labor compensation paid to the goods sector fell from 19.7 percent to 12.5 percent. The fraction of total labor compensation is our preferred measure of a sector's employment share because it maps directly into our model, where we measure labor inputs in terms of effective hours rather than raw hours. As figure 1 shows, this measure moves closely with more common measures such as the share of employees in the goods sector. The construction-sector share of labor compensation rose from 4.4 percent in 1992 to 5.7 percent in 2006, as construction boomed prior to the financial crisis in 2008–2009. Employment in construction then started to fall, and by 2012, the construction-sector share of labor compensation was again 4.4 percent. Reallocation away from goods has, thus far, been permanent, while reallocation into construction was temporary.

3. Model

Our model has two countries: the United States (*us*) and the rest of the world (*rw*). A period is one year. Representative households in each country work, consume, and save to maximize utility subject to a sequence of budget constraints. Competitive firms produce commodities that serve final and intermediate uses both at home and abroad.

Agents in our model have perfect foresight except for in the first period, 1992. We model the saving glut as an unanticipated and temporary — though decades-long — change in the rest of the world's willingness to lend. This assumption captures our view that U.S. households in the early 1990s did not anticipate the large trade deficits their country would run over the next two decades.

3.1 Gross output production

The United States produces goods (*g*), services (*s*), and construction (*c*); the rest of the world produces only goods and services. Gross output is produced according to a nested constant

elasticity of substitution (CES) technology that uses capital, labor, and intermediate inputs of goods and services from both countries. As in the data, construction is not used as an intermediate input; it is only used in the production of investment goods. We use superscripts to index countries or final demand sectors, destination first, and subscripts to index production sectors, destination first. To simplify exposition, we describe only the U.S. economy below. Except where noted, we model the rest of the world analogously.

Gross output in sector i , y_{it}^{us} , is an aggregate of value added, v_{it}^{us} , and an intermediate bundle, m_{it}^{us} ,

$$(1) \quad y_{it}^{us} = \Lambda_i^{us} \left\{ \lambda_i^{us} \left(v_{it}^{us} \right)^\eta + (1 - \lambda_i^{us}) \left(m_{it}^{us} \right)^\eta \right\}^{\frac{1}{\eta}}.$$

The parameter λ_i^{us} governs the share of value added in gross output, and η governs the elasticity of substitution between value added and the intermediate input bundle. Λ_i^{us} is a constant scaling factor used to facilitate calibration.

Value added in sector i is a Cobb-Douglas aggregate of capital, k , and labor, ℓ ,

$$(2) \quad v_{it}^{us} = A_i^{us} \left(k_{it}^{us} \right)^{\alpha_i} \left(\bar{\gamma}_{it}^{us} \ell_{it}^{us} \right)^{1-\alpha_i},$$

where α_i is capital's share of value added, $\bar{\gamma}_{it}^{us}$ is labor productivity, and A_i^{us} is a constant scaling factor. We abstract from capital in the rest of the world,

$$(3) \quad v_{it}^{rw} = A_i^{rw} \bar{\gamma}_{it}^{rw} \ell_{it}^{rw}.$$

Labor productivity is exogenous and grows at different rates in each sector and country. The labor productivity process $(\bar{\gamma}_{it}^{us}, \bar{\gamma}_{it}^{rw})$ is where we introduce our first exogenous driving force: asymmetric productivity growth. We use bars to distinguish labor productivity and other exogenous time-varying parameters from constant parameters and equilibrium variables.

The composite intermediate bundle m_{it}^{us} is composed of intermediate purchases from each traded sector j in each country $(m_{ijt}^{us,us}, m_{ijt}^{us,rw})$,

$$(4) \quad m_{it}^{us} = \Pi_i^{us} \left\{ \sum_{j=g,s} \pi_{ij}^{us} \left[\mu_{ij}^{us} \left(m_{ijt}^{us,us} \right)^{\zeta_j} + (1 - \mu_{ij}^{us}) \left(m_{ijt}^{us,rw} \right)^{\zeta_j} \right]^{\frac{\xi}{\zeta_j}} \right\}^{\frac{1}{\xi}}.$$

Intermediate goods are aggregated in two stages. First, intermediate inputs from each country are aggregated within the goods and services sectors separately. Second, goods and services

intermediates are combined to form the aggregate intermediate bundle used in industry i . Consider, for example, the use of tires and steel in the production of automobiles. Car producers first combine domestic and foreign tires into a tire bundle, and do the same for steel. The tire bundle and steel bundle are then combined to form the total intermediate bundle.

The parameter Π_i^{us} is a constant scaling factor. π_{ij}^{us} governs the share of sector- j intermediates in the aggregate intermediate bundle, and μ_{ij}^{us} govern the domestically sourced share of sector- j intermediates in the sector- i intermediate bundle. ξ governs the elasticity of substitution between intermediates from different sectors, while ζ_j governs the elasticity of substitution between sector- j intermediates from different source countries. This specification allows us to calibrate a production structure in which goods and services are strong complements in intermediate production ($1/(1-\xi) \approx 0$), but an input from the rest of the world is a substitute for an input from the United States ($1/(1-\zeta_j) \geq 1$).

Gross output is produced by perfectly competitive firms. A representative firm in sector i in the United States chooses capital, labor, and intermediate inputs to maximize profits,

$$(5) \quad q_{it}^{us} y_{it}^{us} - w_t^{us} \ell_{it}^{us} - r_{kt}^{us} k_{it}^{us} - \sum_{j=g,s} \left(q_{jt}^{us} m_{ijt}^{us,us} + q_{jt}^{rw} m_{ijt}^{us,rw} \right),$$

subject to nonnegativity constraints and (1), (2), and (4). The wage is w_t^{us} , r_{kt}^{us} is the capital rental rate, and the prices of sector- i gross output sourced from the two countries are q_{it}^{us} and q_{it}^{rw} . Firms in the rest of the world face a similar problem but do not use capital to produce.

3.2 Final demand

The United States has three types of final demand: private consumption (ush), public consumption (usg), and investment (usx). The rest of the world has only consumption. Households in both countries and the U.S. government consume goods and services produced at home and abroad. Investment in the United States is produced using goods and services from home and abroad, and with local construction services.

Similar to the aggregation scheme for intermediate inputs, final demand for output from sector i is a composite of domestic and foreign products. To allow for differences in shares and elasticities across final uses, we construct these composites separately for each category of final

demand. Household consumption of output from tradable sector $i = g, s$ is an aggregate of purchases from home and abroad $(c_{it}^{ush,us}, c_{it}^{ush,rw})$,

$$(6) \quad c_{it}^{ush} = \Theta_i^{ush} \left(\theta_i^{ush} (c_{it}^{ush,us})^{\sigma_i} + (1 - \theta_i^{ush}) (c_{it}^{ush,rw})^{\sigma_i} \right)^{\frac{1}{\sigma_i}}.$$

The parameter σ_i governs the elasticity of substitution in final demand between domestic and foreign products. As with intermediate inputs, we allow goods and services to have different elasticities. The parameter θ_i^{ush} governs the domestically sourced share of household expenditure on sector- i output, and Θ_i^{ush} is a constant scaling factor. The price of sector- i household consumption, p_{it}^{ush} , is given by the standard CES price index.

The government's consumption of sector- i output, c_{it}^{usg} , and the investment producers' use of sector- i output, x_{it}^{usx} , are constructed using analogous versions of (6), and their prices, p_{it}^{usg} and p_{it}^{usx} , are given by the analogous constant elasticity of substitution price indices. The associated share parameters are $(\theta_i^{usg}, \theta_i^{usx})$ and the scale parameters are $(\Theta_i^{usg}, \Theta_i^{usx})$.

3.2.1 Financial assets and exchange rates

Before we can discuss investment production, households, and the government, we need to describe the bond market. The U.S. government, households in the United States, and households in the rest of the world have access to a one-period, internationally traded bond, b_t , that is denominated in units of the U.S. consumer price index (CPI),

$$(7) \quad P_t^{us} = \frac{P_{gt}^{ush} c_{g,1992}^{ush} + P_{st}^{ush} c_{s,1992}^{ush}}{P_{g,1992}^{ush} c_{g,1992}^{ush} + P_{s,1992}^{ush} c_{s,1992}^{ush}}.$$

U.S. households can also save by investing in the U.S. capital stock, but will be indifferent between holding capital and bonds, as they pay the same return in equilibrium. The real interest rate in units of the U.S. CPI is

$$(8) \quad 1 + r_{t+1} = P_t^{us} / Q_t,$$

where Q_t is the bond price. We compute the real exchange rate in our model using consumer price indices, just as we do in the data,

$$(9) \quad rer_t = P_t^{rw} / P_t^{us},$$

where the rest of the world's CPI, P_t^{rw} , is defined analogously to the U.S. CPI in (7).

We model a single bond, but the equilibrium of our baseline model is equivalent to one in which both governments and households issue debt. The deterministic nature of the model gives rise to Ricardian equivalence (except at the unexpected onset of the saving glut), so the split between public and private debt is essentially irrelevant. We have experimented with a stochastic version of the model in which public and private debt are distinct, but we find this change to be quantitatively insignificant.

3.2.2 Investment

Investment is produced by perfectly competitive firms operating a constant returns to scale technology; we focus on the representative firm. Investment is an aggregate of goods and services composites, $(x_{gt}^{usx}, x_{st}^{usx})$, and the output of the U.S. construction sector,

$$(10) \quad x_t^{usx} = E^{usx} \left(\varepsilon_g^{usx} (x_{gt}^{usx})^\nu + \varepsilon_s^{usx} (x_{st}^{usx})^\nu + \varepsilon_c^{usx} (y_{ct}^{us})^\nu \right)^{\frac{1}{\nu}},$$

where E^{usx} is a constant scale factor, ε_i^{usx} governs the share of goods, services, and construction in investment expenditure, and ν governs the elasticity of substitution in investment production. The price of investment, p_t^{usx} , is given by the standard CES price index.

3.2.3 Households

Each country is populated by a continuum of identical households. We draw a distinction between the total (\bar{n}_t^{us}) and working-age ($\bar{\ell}_t^{us}$) populations so that our model captures the impact of demographic changes, both within and across countries, on households' incentives to borrow or save. We normalize the time available for work and leisure in each period to one.

The representative household in the United States chooses consumption of goods and services, investment, bond holdings, and labor supply to maximize lifetime utility,

$$(11) \quad \sum_{t=0}^{\infty} (\beta^{us})^t \frac{1}{\psi} \left\{ \left(\varepsilon_g^{ush} \left(\frac{c_{gt}^{ush} - \bar{c}_{gt}^{ush}}{\bar{n}_t^{us}} \right)^\rho + \varepsilon_s^{ush} \left(\frac{c_{st}^{ush} + \bar{c}_{st}^{ush}}{\bar{n}_t^{us}} \right)^\rho \right)^{\frac{\phi^{us}}{\rho}} \left(\frac{\bar{\ell}_t^{us} - \ell_t^{ush}}{\bar{\ell}_t^{us}} \right)^{1-\phi^{us}} \right\}^\psi,$$

subject to the budget constraints and the law of motion for capital,

$$(12) \quad p_{gt}^{ush} c_{gt}^{ush} + p_{st}^{ush} c_{st}^{ush} + p_t^{usx} x_t^{usx} + Q_t b_{t+1}^{ush} = w_t^{us} \ell_t^{ush} + P_t^{us} b_t^{ush} + (1 - \tau_k^{us}) r_{kt}^{us} k_t^{us} - T_t^{us},$$

$$(13) \quad k_{t+1}^{us} = (1 - \delta)k_t^{us} + x_t^{usx},$$

appropriate nonnegativity constraints, initial conditions for the capital stock and bond holdings, and the no-Ponzi condition. Households pay constant proportional taxes, τ_k^{us} , on capital income and receive a lump-sum tax or transfer, T_t^{us} . The parameter ρ governs the elasticity of substitution between goods and services in household consumption, ϕ^{us} governs the share of time devoted to leisure, $(\varepsilon_g^{ush}, \varepsilon_s^{ush})$ govern the shares of goods and services in household consumption, and ψ determines the intertemporal elasticity of substitution.

Following Buera and Kaboski (2009) and Herrendorf et al. (2013), we use Stone-Geary preferences to model nonhomotheticities that generate structural change. The subsistence requirement for goods, \bar{c}_{gt}^{ush} , and the endowment of services, \bar{c}_{st}^{ush} , vary over time to ensure that the model retains consistency with balanced growth.¹

The rest of the world's households solve a similar but simpler problem. They choose consumption of goods and services, labor supply, and bond holdings to maximize lifetime utility,

$$(14) \quad \sum_{t=0}^{\infty} (\beta^{rw})^t \bar{\omega}_t^{rw} \frac{1}{\psi} \left\{ \left(\varepsilon_g^{rw} \left(\frac{c_{gt}^{rw} - \bar{c}_{gt}^{rw}}{\bar{n}_t^{rw}} \right)^{\rho} + \varepsilon_s^{rw} \left(\frac{c_{st}^{rw} + \bar{c}_{st}^{rw}}{\bar{n}_t^{rw}} \right)^{\rho} \right)^{\frac{\phi^{rw}}{\rho}} \left(\frac{\bar{\ell}_t^{rw} - \ell_t^{rw}}{\bar{\ell}_t^{rw}} \right)^{1-\phi^{rw}} \right\}^{\psi},$$

subject to the budget constraints

$$(15) \quad p_{gt}^{rw} c_{gt}^{rw} + p_{st}^{rw} c_{st}^{rw} + Q_t b_{t+1}^{rw} = w_t^{rw} \ell_t^{rw} + P_t^{us} b_t^{rw},$$

nonnegativity constraints, and a no-Ponzi condition. The parameter $\bar{\omega}_t^{rw}$, which shifts the rest of the world's intertemporal demand, is the second exogenous driving force in our model. We will calibrate the series $\{\bar{\omega}_t^{rw}\}_{t=1992}^{2012}$ so that the U.S. trade balance in our model's saving-glut equilibrium matches the data. After 2012, $\bar{\omega}_t^{rw}$ gradually reverts to one,

$$(16) \quad \bar{\omega}_{t+1}^{rw} = \varphi \bar{\omega}_t^{rw} + (1 - \varphi).$$

¹ Kongsamut et al. (2001) show that, in a closed-economy model with Stone-Geary preferences, one can ensure a balanced growth path by choosing the subsistence requirements so that they cancel out in the household's budget constraint. In closed-economy models, balanced growth paths are unique. Our model has a continuum of balanced growth paths with different relative prices, so there are no subsistence parameters that generically satisfy this knife-edge condition. Boppart (2014) constructs a closed-economy model that generates balanced growth without this parameter restriction. We leave extending that model to an open economy with unbalanced trade to future research.

3.2.4 U.S. government

The government in the United States levies taxes and sells bonds to finance exogenously required consumption expenditures. The budget constraint is

$$(17) \quad p_{gt}^{usg} c_{gt}^{usg} + p_{st}^{usg} c_{st}^{usg} + Q_t b_{t+1}^{usg} = \tau_k^{us} r_{kt}^{us} k_t^{us} + T_t^{us} + P_t^{us} b_t^{usg}.$$

Total government consumption expenditures and debt as fractions of GDP follow exogenous time paths, \bar{c}_t^{usg} and \bar{b}_t^{usg} . These paths follow the historical data for 1992–2012, and in the subsequent years, the paths follow the Congressional Budget Office (CBO) projections (Congressional Budget Office, 2012). We allow the lump-sum tax, T_t^{us} , to vary as necessary to ensure that the government’s budget constraint is satisfied. The government’s objective in each period is to choose goods and services consumption to maximize a CES “utility function” subject to the constraint that total consumption expenditures are at the required level,

$$(18) \quad \max_{c_{gt}^{usg}, c_{st}^{usg}} \left(\varepsilon_g^{usg} (c_{gt}^{usg})^\nu + \varepsilon_s^{usg} (c_{st}^{usg})^\nu \right)^{\frac{1}{\nu}}$$

subject to $p_{gt}^{usg} c_{gt}^{usg} + p_{st}^{usg} c_{st}^{usg} = \bar{c}_t^{usg} GDP_t^{us}$,

where $(\varepsilon_g^{usg}, \varepsilon_s^{usg})$ govern the shares of goods and services in government consumption and ν governs the elasticity of goods and services in government consumption. Government debt is simply given by $b_t^{usg} = \bar{b}_t^{usg} GDP_t^{us}$.

The lump-sum tax implies that our model exhibits near-Ricardian equivalence, which means that the timing of taxes and government borrowing is almost irrelevant. Ricardian equivalence breaks down only when we introduce unexpected events such as the saving glut. In our sensitivity analysis, however, we show that deviations from Ricardian equivalence have little impact on our results.

3.3 Market clearing

Market clearing for gross output of goods and services in each country $l = \{us, rw\}$ requires that gross output, y_{it}^l , equals the sum of all intermediate and final demand (recall that we do not model the government or the investment sector in the rest of the world),

$$(19) \quad y_{it}^l = \sum_{k=us, rw} \sum_{j=g, s} \left(m_{ijt}^{us, k} + m_{ijt}^{rw, k} + c_{jt}^{ush, k} + c_{jt}^{rw, k} + c_{jt}^{usg, k} + x_{jt}^{usx, k} \right), \quad i = g, s.$$

Construction market clearing is trivial since construction is used only to produce investment. Market clearing in factor markets requires

$$(20) \quad \ell_t^{us} = \sum_{i=g,s,c} \ell_{it}^{us}, \quad \ell_t^{rw} = \sum_{i=g,s} \ell_{it}^{rw}, \quad k_t^{us} = \sum_{i=g,s,c} k_{it}^{us},$$

and finally, bond markets must clear,

$$(21) \quad b_t^{ush} + b_t^{usg} + b_t^{rw} = 0.$$

3.4 Equilibrium and balanced growth paths

An *equilibrium* in our model, given a sequence of time-varying parameters $\{\bar{\omega}_t^{rw}, \bar{c}_t^{usg}, \bar{b}_t^{usg}, (\bar{n}_t^l, \bar{\ell}_t^l)_{l \in \{us, rw\}}, (\bar{\gamma}_{it}^l)_{l \in \{us, rw\}, i \in \{g, s, c\}}\}_{t=0}^{\infty}$ and initial conditions $(b_{1992}^{ush}, b_{1992}^{usg}, k_{1992}^{us})$, is sequences of model variables such that: households in the United States and the rest of the world maximize utility subject to their constraint constraints; prices and quantities satisfy marginal product pricing conditions for all commodities; prices and quantities are such that all production activities earn zero profits; all commodity, factor, and bond market clearing conditions are satisfied; and the U.S. government solves its consumption allocation problem.

When $\{\bar{\omega}_t^{rw}, \bar{c}_t^{usg}, \bar{b}_t^{usg}, (\bar{n}_t^l, \bar{\ell}_t^l)_{l \in \{us, rw\}}\}_{t=0}^{\infty}$ are constant, and labor productivity in each sector grows at the same constant rate, $\{\bar{\gamma}_{it}^l / \bar{\gamma}_{i,t-1}^l = g_\gamma\}_{t=0}^{\infty}$ for $l \in \{us, rw\}$ and $i \in \{g, s, c\}$, the model converges to a balanced growth path. A *balanced growth path* in our model is an equilibrium in which all quantities grow at the constant rate of productivity growth, $g_\gamma - 1$ (except for labor supply, which is constant), and all relative prices are constant. Since we allow for unbalanced trade, our model has a continuum of balanced growth paths indexed by U.S. net foreign asset levels. The balanced growth path to which a particular equilibrium converges is determined by the equilibrium's initial conditions and time-varying parameters. Consequentially, solving the model is not merely transiting between known initial conditions and a known balanced growth path. We must solve for the transition and the balanced growth path simultaneously. This complexity influences our calibration approach, as we describe in the next section.

To solve the model, we require that the model converges to a balanced growth path by 2092, 100 years after our starting date. This implies that sectoral productivity growth rates must eventually be equal, the other exogenous time-varying parameters must be constant, and household preferences must be homothetic. We do not want these end point conditions to distort the behavior

of the model during our period of interest (1992–2012). To achieve this, we allow the time-varying parameters to take their calibrated values during 1992–2042. For the 50 years that follow, sectoral productivity growth rates converge to equality, the other exogenous time-varying parameters converge to constants, and the subsistence consumption requirements converge to zero. Our results during 1992–2012 are not sensitive to increasing the time it takes to converge to the balanced growth path beyond 100 years.

We construct two equilibria in our model: a *benchmark* in which the saving glut is followed by a gradual rebalancing of U.S. trade and a *no-saving-glut counterfactual* in which the saving glut does not occur and U.S. trade is roughly balanced forever. The differences between these two equilibria are completely due to the paths of $\{\bar{\omega}_t^{rw}, \bar{c}_t^{usg}, \bar{b}_t^{usg}\}_{t=0}^{\infty}$; all other parameters and the initial conditions do not change.

4. Calibration

We calibrate the model so that the counterfactual equilibrium (the one without the saving glut) replicates the U.S. data in 1992. Our assumption is that the saving glut was not foreseen in 1992. We view this as the most natural exercise, but our results are not sensitive to calibrating a benchmark equilibrium in which agents foresee the saving glut.

We take the elasticities of substitution from the relevant literature and calibrate the remainder of the model’s production and preference parameters so that the equilibrium replicates the 1992 world input-output matrix (table 1) we have constructed from the World Input-output Database (WIOD). In the next section, we show that, while we have not calibrated the elasticities, the model does a good job in matching observed changes in several key expenditure shares — evidence that our chosen elasticities are appropriate. Table 2 summarizes the parameter values. More details are provided in the online supplementary materials available at www.econ.umn.edu/~tkehoe.

4.1 Gross output parameters

We choose the scaling factors Λ_i^l , A_i^l , and Π_i^l ($l = us, rw$) so that U.S. GDP in 1992 is 100 and all gross output prices in 1992 are one. Following Atalay (2014), we set the elasticity of substitution between value added and intermediate inputs, $1/(1-\eta)$, to 0.05, and the elasticity of substitution between intermediate inputs from different sectors, $1/(1-\xi)$, to 0.03. For the within-sector elasticities of substitution between intermediates from different countries, $1/(1-\zeta_i)$, we use three for goods and one for services. These choices are near the upper and lower bounds of the

aggregate Armington elasticities used in the literature. Given these values, we choose the share parameters, $(\pi_{ij}^{us}, \pi_{ij}^{rw})$ and $(\mu_{ij}^{us}, \mu_{ij}^{rw})$, so that the equilibrium replicates the quantities in our world input-output matrix.

We calibrate the U.S. labor productivity growth parameters, $\bar{\gamma}_{it}^{us}$, so that real value added per worker in the model grows at the same rate as in the data for each sector. Using data on value added and labor compensation from the Bureau of Economic Analysis (BEA) for 1992–2012, we find that the average annual growth rates of real value added per effective worker are 4.4 percent in goods, 1.3 percent in services, and -0.84 percent in construction. In the model’s first 50 years (1992–2042), we set $\bar{\gamma}_{it}^{us}$ so that real value added per worker in the model grows at the same rates as in the data. Starting in 2042, all of the sectoral productivity growth rates begin to converge to 2 percent per year to ensure that the equilibrium converges to a balanced growth path.² This convergence process takes 25 years. Our results, which focus on the period 1992–2012, are not sensitive to this timing assumption. For the rest of the world, we calculate growth rates of $\bar{\gamma}_{it}^{rw}$ so that the price of goods relative to services follows the same path in the rest of the world as in the United States.

4.2 Final demand parameters

Again, we choose scaling factors Θ_i^l so that 1992 final demand prices are one. The WIOD data do not distinguish between public and private consumption, so we do not have separate data on changes in public and private consumption expenditure shares. We therefore set both $1/(1-\rho)$ and $1/(1-\nu)$ to the estimate in Atalay (2014) of 0.65 — about halfway between the commonly used value of 0.5 and the estimate of 0.85 by Herrendorf et al. (2013). We set the elasticity of substitution between different commodities in investment, $1/(1-\nu)$, to one (Bems, 2008). The final demand Armington elasticities, $1/(1-\sigma_i)$, are two for goods and one for services. As with

² In Ngai and Pissarides (2007), when sectors differ permanently in productivity growth, long-run balanced growth is possible if and only if the elasticity of substitution across sectors in intermediate use is one. Since we use a low elasticity of substitution between intermediate goods and services, we assume convergence across sectors in productivity growth rates to ensure balanced growth. This assumption has no significant impact on the equilibrium dynamics during our period of interest, 1992–2012.

the gross output production structure, given our choice of elasticities, we choose values for the share parameters, ε_i^l , so that the equilibrium replicates the world input-output matrix.

4.3 Household parameters

We set the share parameters, ϕ^l , so that households in both countries use one-third of their time for work. We set the U.S. and world discount factors so that the real interest rate on the balanced growth path is 3 percent. Estimates of the real interest rate from academics and government research groups range from 0.5 to 5 percent (McGrattan and Prescott 2000, 2005; King and Low 2014; Congressional Budget Office, 2012, 2015). Given these considerations, we chose the CBO's 2012 projection of 3 percent because it is in line with estimates from the literature, and is consistent with the CBO's budget projections, which we use in modeling government spending and debt. Our main results are not sensitive to reasonable interest rate values.

For the demographic time series \bar{n}_i^l and $\bar{\ell}_i^l$, we use historical data and future projections from the *World Population Prospects: 2010 Revision* (United Nations, 2011). Consistent with our measure of the real exchange rate in Section 2, we construct the rest of the world's population as a weighted average of the United States' 20 largest trade partners. To ensure convergence to a balanced growth path, we assume that populations in the United States and the rest of the world start to converge to constant levels in 2042 and finish converging 25 years later.

The subsistence consumption levels, \bar{c}_i^{ush} , are constant during 1992–2042. Herrendorf et al. (2013) report subsistence consumption levels for 1947 and 2010. We linearly interpolate between these dates to find the 1992 subsistence parameters for the United States, and we assume that, in per capita terms, the rest of the world's subsistence levels are the same as those in the United States: $\bar{c}_{i,1992}^{rw} = (\bar{\ell}_{1992}^{rw} / \bar{\ell}_{1992}^{us}) \bar{c}_{i,1992}^{ush}$. To ensure a balanced growth path in our model, preferences must be homothetic in the long run, so we assume the subsistence parameters converge to zero during 2042–2067. We provide more details in the online supplementary materials.

4.4 Government spending and debt

We calibrate the model under the assumption that agents did not foresee the saving glut. It is likely that the observed path of government spending in the data has been influenced by the saving glut, so we must specify agents' beliefs over future government actions in the absence of the saving glut. In the no-saving-glut counterfactual, we assume that government spending remains at its 1992

level (as a share of GDP) forever and that government debt converges from 48 percent of GDP (the observed level in 1992) to 60 percent of GDP, consistent with the 1992 CBO projections. We set the initial value of household bond holdings, b_{1992}^{ush} , so that total net foreign assets, $b_{1992}^{ush} + b_{1992}^{usg}$, are -7.8 percent of GDP (Lane and Milesi-Ferretti, 2007). When we solve the model with the saving glut, the behavior of government spending and debt will differ as discussed below.

4.5 Capital stock, capital tax, and depreciation

Using the method of Backus et al. (2007), we compute U.S. capital stocks for 1992–2012. To calibrate the depreciation rate, we use the capital stock data, and the consumption of fixed capital, reported by the BEA, to calculate an average annual depreciation rate of 5.3 percent. We set the U.S. capital income tax, τ_k^{us} , to 39 percent, the statutory rate reported by Devereux et al. (2002). We set the initial capital stock, k_{1992} , to 278 percent of GDP, as it is in the data.

5. Quantitative results

Having calibrated the model, we now present our quantitative results. We start with a comparison of our benchmark saving-glut equilibrium to the data, and then examine the no-saving-glut counterfactual to answer the question: what would have happened to the U.S. economy — and what would happen in the future — but for the saving glut?

5.1 The benchmark saving-glut equilibrium

We have calibrated the model to the no-saving-glut counterfactual — the equilibrium in which the saving glut did not occur. We now use the calibrated model to study the impact of the saving glut. In this equilibrium, we keep all of the model parameters as they are in the counterfactual, except for three: We change $\bar{\omega}_t^{rw}$ to create the saving glut, and we change government spending and debt paths to match the observed data. We calibrate $\bar{\omega}_t^{rw}$ so that the equilibrium replicates the aggregate U.S. trade balance during 1993–2012; in the subsequent years, $\bar{\omega}_t^{rw}$ reverts to one according to (16). U.S. government consumption and debt (as a share of GDP) match the data for 1992–2012 and follow projections from the CBO’s 2012 *Long-Term Budget Outlook*, which we adjust in later years to be consistent with balanced growth. More details on our construction of projected government spending and debt are available in the online supplementary materials. Our results are

not sensitive to assumptions about the behavior of government spending and debt, as well as households' expectations about these variables.³

To compute the equilibrium in the model with the saving glut, we use 1993 as the initial period and use the 1993 equilibrium values of capital and bond holdings from the no-saving-glut counterfactual as initial conditions. This specification embodies the assumption that agents in 1992 did not foresee the saving glut and were surprised by the increase in demand for saving in the rest of the world, as well as by the increase in domestic government spending and debt that accompanied it. Once agents learn of the saving glut, they have perfect foresight.

5.1.1 Goodness of fit: Changes in expenditure shares

Before turning to our three key facts, we compare the model with the data along several key dimensions that were not targets of the calibration. These moments are useful in judging the elasticities of substitution that we have chosen from the literature. We focus on changes in U.S. expenditure shares between 1995 and 2011, which map directly to the elasticities of substitution. We report the changes in expenditure shares in the data and the model in table 3.

The elasticity of substitution between value added and intermediates is an important determinant in the change in the intermediate-input share of gross output. We use the estimate from Atalay (2014), $1/(1-\eta) = 0.05$, which makes this share sensitive to changes in the relative price of intermediates. As we report in table 3, the change in the intermediate-input share of gross output between 1995 and 2011 is -1.04 percent in the data and -1.50 percent in the model. The elasticity of substitution between intermediate inputs of goods and services affects the sensitivity of the goods share of total intermediate use to the relative price of goods; we use the estimate from Atalay (2014), $1/(1-\xi) = 0.03$. In the data, the goods share of total intermediate use falls by 6.8 percent, compared with 5.4 percent in the model.

The elasticities of substitution between goods and services in public and private consumption are $1/(1-\rho) = 0.65$ and $1/(1-\nu) = 0.65$. As shown in Herrendorf et al. (2013), these elasticities must be consistent with the underlying model; elasticities in models of gross output are significantly larger than those in models of value added. We use the estimate in Atalay (2014),

³ For example: If we assume government spending and debt are constant, the saving glut is responsible for 17 percent of the decline in goods-sector employment and rises to 20 percent if we assume households perfectly foresee the actual path of government spending and debt during the saving glut.

which is derived from a gross output model with an input-output structure and distinct elasticities of substitution between sectoral output in final and intermediate uses. To evaluate the appropriateness of our elasticity, we examine the change in the goods share of aggregate consumption. In the data, this share falls by 1.7 percent, compared with 2.7 percent in the model.

The elasticity of substitution between goods and services in investment is $1/(1-\nu)=1$, as found in Bems (2008). The change in the goods share of investment in the data is -0.2 percent, compared with the unchanged share in our model.

Lastly, the elasticities between domestic and foreign products in intermediate and final uses, $1/(1-\zeta_i)$ and $1/(1-\sigma_i)$, determine the volatility of the sectoral trade balances. The standard deviation of the goods trade balance is 1.5 in the data, compared with 1.5 in the model. The lower elasticity of substitution between services produced at home and abroad generates less volatility in the services trade balance, which matches the data well. The standard deviation of the services trade balance is 0.24 in the data, compared with 0.20 in the model.

5.1.2 Goodness of fit: Replicating the three key facts

In figure 5, we plot the real exchange rate in the baseline model against the data and the no-saving-glut counterfactual. Our model does a good job in matching the magnitude of the appreciation during 1992–2012: The real exchange rate appreciates by 27.9 percent in the data and 27.1 percent in the model before beginning to depreciate. The baseline model, however, fails to capture the timing of the depreciation. In the data, the real exchange rate begins to depreciate in 2002, four years before the trade deficit begins to shrink. In our model, the real exchange rate moves in tandem with the trade balance; it does not begin to depreciate until 2006.

The lag between the real exchange rate and trade balance movements in the data, often referred to as the “J-curve” (Backus, et al., 1994), is difficult to generate in models of aggregate trade such as ours. Alessandria and Choi (2015), for example, show that heterogeneous firms that price to market and make dynamic decisions about entering and exiting the export market are crucial elements in generating a realistic J-curve. The sort of frictions that generate J-curves at the business cycle frequency, however, are unlikely to generate the sort of four-year lag that we see in the U.S. data. To incorporate this lag into the model in Section 6, we mechanically introduce wedges so that the model matches both the trade balance and real exchange rate during 1992–2012, and show that this has little impact on our model’s predictions for sectoral labor reallocation.

In figure 6, we plot the sector-level trade balances in the model and data. The model closely matches the trade balances for both goods and services between 1992 and 2012. This aspect of the model's performance is due in part to our choice of Armington elasticities. Table 3 shows that in both the model and the data, the goods trade balance is more volatile than the services trade balance. Had we used the same elasticities in both sectors, the goods trade balance would not have moved enough, while the services trade balance would have been too volatile.

In figure 7, we plot the employment shares for the goods and construction sectors in the model and the data. The baseline model captures 99.2 percent of the decline in the goods-sector employment share between 1992 and 2012. Our model's ability to capture the full extent of the reallocation of labor away from the goods sector makes it well suited to answering the counterfactual question at the heart of the paper: What would have happened to the goods-sector employment share in the absence of the saving glut?

Figure 7 also shows that the model captures aspects of the construction-sector employment share between 1992 and 2012. During 1993–2006, the construction-sector employment share rises in both the model and the data, although our model generates a larger increase. The subsequent bust is smaller in the model than in the data, primarily because we have not introduced the financial crisis of 2008–2009 in any form other than the way in which it affected the trade balance. If we were to introduce a more accurate model of the crisis, we would undoubtedly do better in this regard, but this is not the focus of our study.

5.2 The impact of the saving glut

We now use the calibrated model to study the impact of the saving glut by comparing the baseline saving-glut equilibrium with the no-saving-glut counterfactual. This comparison is informative about the impact of the saving glut in both the past and the future.

Our first question concerns the trade balance: What would the U.S. trade balance have been had the rest of the world's effective discount factor remained constant? Figure 4 shows that, in the absence of the saving glut, U.S. trade would have been roughly balanced throughout the 1992–2012 period; the entire cumulative U.S. trade deficit during this period is due to increased demand for saving in the rest of the world. This is not an obvious or a trivial result — initial conditions and differences in demographic and industrial structures across countries, coupled with asymmetric productivity growth across sectors, can generate cross-country differences in saving behavior. Our results tell us that these factors are not important in explaining the behavior of the U.S. trade

balance over the past two decades. Looking into the future, our model predicts that the trade balance will be about 1 percent of GDP higher by 2024 than it would have been if the saving glut had not occurred. This is driven by the repayment of the debt incurred by the United States during the saving glut.

Examining the counterfactual sector-level trade balances in figure 6, we see that the saving glut's impact on the trade balance has been concentrated in the goods sector; the services trade balance would have been similar had the saving glut not occurred. This is also true in the future. The trade surpluses the United States will run to pay back its saving-glut debt will come almost entirely from the goods sector. The United States will run a surplus in services as well, but by 2024 this surplus would be the same as if the saving glut had not occurred.

Our counterfactual real exchange rate mirrors the aggregate trade balance. As evident in figure 5, in the counterfactual equilibrium, the real exchange rate is roughly unchanged between 1992 and 2014. Comparing the counterfactual with the baseline, we see that all of the real exchange rate appreciation between 1992 and the mid-2000s is attributable to the saving glut, as is the subsequent depreciation. In the future, as the United States repays its debt to the rest of the world, its real exchange rate will depreciate above its path in the no-saving-glut counterfactual.

Our primary interest is the impact of the saving glut on labor reallocation in the United States: What portion of the decline in the goods-sector employment share during 1992–2012 is attributable to the saving glut? To answer this question, we compare the drop in the goods-sector employment share in the counterfactual equilibrium with the drop in the baseline equilibrium. We measure both in the same units: the fraction of the decline in the goods-sector employment share in the data. The decline in the no-saving-glut counterfactual is 84.1 percent of the observed decline, and the decline in the baseline equilibrium is 99.2 percent. The difference between the two, 15.1 percent, is the decrease in the goods-sector employment share that is driven by the saving glut. Looking to the future, figure 7 shows that the reversal of the saving glut will have a negligible impact on goods-sector employment. The share of employment in the goods sector would be about the same in 2024 if the saving glut did not happen; the surplus in goods trade necessary to repay the debt will have little impact on the allocation of labor across sectors. Eliminating the U.S. trade deficit will not bring back employment to goods production.

Our focus is on the saving glut's impact on structural change, but we can also measure the impact of traditional structural change forces in our model: differential productivity growth and

nonhomothetic preferences. When we set the subsistence consumption levels to zero, so that preferences are homothetic, the benchmark model accounts for 93.5 percent of the decline in the goods-sector employment share, compared with 99.2 percent in the baseline model; differential productivity growth is clearly the dominant driver of structural change in our model. Our assumption about preferences does not change our finding regarding the saving glut. The saving glut is responsible for 14.4 percent of the decline in the goods-sector employment share in the homothetic model, compared with 15.1 percent in the baseline model.

6. Sensitivity analysis

We have conducted a wide range of additional experiments with our model, and we find that our main result — the modest impact of the saving glut on U.S. structural change — is robust. We limit our discussion here to only three sets of sensitivity analyses. The first modifies the model to improve its ability to match the timing of the U.S. real exchange rate. In the second, we calibrate our model to counterfactual input-output tables to shed light on the role of intermediate input linkages in driving structural change. In the last analysis, we model trade deficits as driven by domestic factors rather than increased demand for saving in the rest of the world. We summarize the results of our sensitivity analysis in table 4. In the supplementary materials available online, we demonstrate the robustness of our results to the household’s labor-leisure choice.

6.1 Trade wedges

Figure 5 shows that our model matches the magnitude of the real exchange rate appreciation and depreciation in the data, but misses on the timing. The real exchange rate begins to appreciate in 2002, but the trade deficit does not start to shrink until 2006. To understand the importance of this phenomenon for our results, we study a version of our model with “trade wedges” as in Alessandria et al. (2013). In this model, shipments of goods and services from the United States to the rest of the world are distorted by the time-varying factor $(1 + \tau_{mt}^{rw})$. Taking the trade wedges, $\tau_{m,t}^{rw}$ as given, producers in the rest of the world maximize profits,

$$(22) \quad q_{it}^{rw} y_{it}^{rw} - w_t^{rw} \ell_{it}^{rw} - \sum_{j=g,s} \left((1 + \tau_{m,t}^{rw}) q_{jt}^{us} m_{ijt}^{rw,us} + q_{jt}^{rw} m_{ijt}^{rw,rw} \right).$$

This implies that the prices of goods and services in the rest of the world are now functions of the trade wedges, distorting producers’ and consumers’ demand for products from the United States. We calibrate these wedges so that the model matches exactly both the trade balance and the real exchange rate data during 1992–2012.

In figure 8 we plot the goods-sector employment share in the baseline model and the trade wedge model. The two models behave almost identically between 1992 and 2002; the trade wedge model generates less movement in the goods-employment share between 2002 and 2016, and the two models are almost identical thereafter. The benchmark trade wedge model generates 94.6 percent of the observed decline in the goods-sector employment share between 1992 and 2012, versus the baseline model's 99.2 percent. The no-saving-glut counterfactual is the same in both models. Compared with the baseline model, the trade wedge model attributes a smaller portion of the decline in the goods-sector employment share during 1992–2012 to the saving glut, 10.6 percent compared to 15.1 percent.

6.2 The input-output structure

Input-output linkages quantitatively affect U.S. structural change in two ways. First, services are stronger complements in intermediate use than in final use, so expenditures on intermediate goods fall more than expenditures on final goods in response to asymmetric productivity growth. This is evident in table 3: the goods share of U.S. intermediates fell more than the goods share of U.S. consumption between 1995 and 2011, -6.8 percent versus -1.7 percent. Second, compared with services, goods are used more intensively as intermediates. The ratio of gross output to value added is 2.9 for goods and 1.7 for services. We construct two counterfactual versions of our input-output table to illustrate the role of these two mechanisms.

In our first counterfactual, we set all of the intermediate input requirements to zero while keeping sectoral value added fixed. This exercise captures the full extent of the input-output linkages' contribution to U.S. structural change. In this production structure, sectoral output is now value added rather than gross output, so we must adjust the elasticity of substitution between goods and services in consumption (Herrendorf et al. 2013). We set the elasticity of substitution between goods and services in consumption to 0.03, the same as our elasticity of substitution between goods and services in intermediate use. This is very close to the suggested elasticity of 0.02 from Herrendorf et al. (2013). In this version of the model, the benchmark equilibrium accounts for 66.2 percent of the decline in the goods-sector employment share. This is substantially less than the baseline model's 99.2 percent, indicating that there remains an important role for input-output linkages in driving U.S. structural change.

Our second counterfactual input-output table illustrates how heterogeneity in intermediate use across sectors affects structural change. Here we retain the input-output structure, but equalize

the gross output to value added ratios in all U.S. sectors to 2.0, the aggregate gross output to value added ratio. Our results suggest that the intensive use of goods as intermediates has hastened the decline in goods-sector employment. This version of the model accounts for 86.1 percent of the observed decline in the goods-sector employment share, 13.0 percent less than the baseline model. The intuition is that when goods are used less intensively as intermediate inputs, there is less room for expenditures on intermediate goods to fall in response to relative price changes driven by asymmetric productivity growth.

In all of our experiments with counterfactual input-output matrices, the saving glut accounts for a similar portion of the decline in the goods-sector employment share as in the baseline model. Although the input-output structure of the U.S. economy is important in explaining structural change, it does not play an important role in the impact of trade deficits on structural change.

6.3 Global saving glut or domestic saving drought?

We have adopted the global-saving-glut hypothesis proposed by Bernanke (2005), which posits that U.S. international borrowing since the 1990s has been driven primarily by increased demand for saving in the rest of the world. A number of authors (Chinn and Ito, 2007; Gruber and Kamin, 2007; Obstfeld and Rogoff, 2009) argue that domestic factors such as monetary policy, housing market policy, and innovations in financial markets were the primary causes of U.S. borrowing. To evaluate the merits of these hypotheses, we study a version of our model in which the preferences of U.S. households, rather than households in the rest of the world, drive the U.S. trade balance. Following Chinn and Ito (2007), we call this the *domestic-saving-drought* model. In the saving-drought model, the preferences of U.S. households take the same form as (14), and we calibrate the U.S. preference parameters, $\bar{\omega}_t^{us}$, to match the U.S. trade balance.

To assess which of the models is consistent with the data, we focus on investment. Figure 9 shows that, before the financial crisis of 2009, U.S. investment as a fraction of GDP rose steadily. This is consistent with the saving-glut hypothesis: U.S. households took advantage of cheap foreign goods to increase both investment and consumption, since the relative value they placed on future consumption remained unchanged. If U.S. borrowing was instead driven by reduced domestic demand for saving, U.S. households should have reduced investment in favor of consumption. Except for 1993, the investment rate in the baseline model moves in the same direction as the data. By contrast, beginning in 1997, the investment rate in the saving-drought model falls dramatically, while it continues to rise in the data (except during the 2001 recession,

which we have not incorporated into our model). During the financial crisis of 2008–2009, which we have not explicitly incorporated into our model, the investment rate falls in the baseline model and the data, but rises in the saving-drought model. Overall, the correlation between the saving-glut model’s investment rate and the data in first differences is 0.87; the same correlation for the saving-drought model is -0.56 .

7. Concluding remarks and directions for future research

We have developed a model of the United States and the rest of the world with two exogenous driving forces: increased foreign demand for saving — the saving glut — and faster productivity growth in goods compared with services and construction. The model accounts for three key facts about the U.S. economy during 1992–2012: (1) the real exchange rate appreciated, then depreciated; (2) the trade balance dynamics are driven almost entirely by the goods trade balance; and (3) labor shifted away from the goods sector toward services and construction. We have used our model to show that, while faster productivity growth in the goods sector is responsible for the bulk of the long-run shift in employment away from that sector, the saving glut hastened this change during 1992–2012.

Although the saving glut’s impact on goods-sector employment is temporary, this does not imply that the saving glut has not had an impact on the U.S. economy: The U.S. economy’s current long-run trajectory is very different from the one it would have taken had the saving glut not occurred. Figures 4 and 5 illustrate this point by plotting the aggregate trade balance and real exchange rate in our benchmark saving-glut equilibrium against the no-saving-glut counterfactual. In the counterfactual, U.S. trade is approximately balanced in the long run, since the United States has little debt to repay. Because the saving glut did happen, however, our model predicts that the United States will have to run a trade surplus of around 1 percent of GDP in perpetuity. To do so, the U.S. real exchange rate will remain permanently depreciated by about 6 percent compared with its path in the no-saving-glut counterfactual.

Our analysis identifies two puzzles. Here we discuss these puzzles and point out directions that future research could take in addressing them.

The first puzzle is: Why did U.S. borrowing continue to increase once the real exchange rate began to depreciate after 2002? In other words, why did U.S. purchases of foreign goods and services continue to increase once foreign goods and services started getting more expensive (figure 3)? A partial resolution to this puzzle might be found in the J-curve literature, in that time-

to-build and import pattern adjustment frictions can delay quantities adjusting to price changes. This mechanism, however, is not likely to explain the substantial four-year lag. Another, perhaps more plausible, explanation is the increase in China's importance in U.S. borrowing during the period. In figure 10 we decompose the U.S. real exchange rate into two components: (1) the bilateral real exchange rate with China; and (2) the real exchange rate with the United States' other major trade partners. We see that the overall real exchange rate and the exchange rate with non-China countries move closely in the early part of the period, but diverge in the latter part. Following 2002, the aggregate real exchange rate behaves more like the real exchange rate with China. To accurately capture this, we would need to model an economy with (at least) three countries and some sort of asset market segmentation, where countries like China choose to lend to the United States rather than to other countries.

The second puzzle is that, in contrast to Bernanke's (2005) judgment, the saving glut has had only a small effect on U.S. real interest rates in the model (see figure 11). The largest difference between the interest rate in the benchmark saving-glut equilibrium and the no-saving-glut counterfactual is 50 basis points. This is consistent with the findings of Warnock and Warnock (2009) and Krishnamurthy and Vissing-Jorgensen (2007), who have estimated that foreign lending drove interest rates down by about 80 basis points over this period.

In our model, the impact of the saving glut on interest rates depends on how substitutable foreign goods are for U.S. goods. With the elasticities that we have chosen, we find that the saving glut generates the right magnitude of the U.S. real exchange rate appreciation, as figure 11 shows, but not the right magnitude in the drop of the U.S. real interest rate. If we make foreign goods more substitutable for U.S. goods, we can generate a larger drop in the U.S. real interest rate during 2006–2012 — although still not as large as the drop observed in the data — but the model would then predict a much smaller appreciation in the U.S. real exchange rate.

To account for the low U.S. real interest rates in the data, we need to look elsewhere, perhaps to the U.S. policies discussed by Obstfeld and Rogoff (2009) and Bernanke et al. (2011). It is worth pointing out, however, that modeling the source of the global imbalances over 1992–2012 as being generated by U.S. saving behavior does not work well. The domestic-saving-drought model we studied in Section 6.4 is successful in generating lower U.S. real interest rates during 1993–2006, as the dollar appreciates, but it generates higher U.S. interest rates during 2006–2012, as the dollar

depreciates. The low interest rates during the entire period pose a puzzle for theories of the saving glut.

8. References

- Acemoglu, D., and V. Guerrieri (2008), “Capital Deepening and Nonbalanced Economic Growth,” *Journal of Political Economy*, 116, 467–498.
- Alessandria, G. and H. Choi (2015), “The Dynamics of the Trade Balance and the Real Exchange Rate: The J-curve and Trade Costs?” unpublished manuscript.
- Alessandria, G., J. Kaboski, and V. Midrigan (2013), “Trade Wedges, Inventories, and International Business Cycles,” *Journal of Monetary Economics*, 60, 1–20.
- Atalay, E. (2014), “How Important are Sectoral Shocks?” unpublished manuscript.
- Backus, D., E. Henriksen, F. Lambert, and C. Telmer (2006), “Current Account Fact and Fiction,” unpublished manuscript.
- Backus, D., E. Henriksen, and K. Storesletten (2007), “Taxes and the Global Allocation of Capital,” *Journal of Monetary Economics*, 55, 48–61.
- Backus, D., P. Kehoe, and F. Kydland (1994), “Dynamics of the Trade Balance and the Terms of Trade: The J-Curve?” *American Economic Review*, 84, 84–103.
- Bems, R. (2008), “Aggregate Investment Expenditures on Tradable and Nontradable Goods,” *Review of Economic Dynamics*, 11, 852–883.
- Bems, R. and K. Jönsson (2006), “Trade Deficits in the Baltic States: How Long Will the Party Last?” *Review of Economic Dynamics*, 9, 179–209.
- Bernanke, B. S. (2005), “The Global Saving Glut and the U.S. Current Account Deficit,” speech at the Sandridge Lecture, Virginia Association of Economists, Richmond, VA, March 10.
- Bernanke, B. S., C. Bertaut, L. P. DeMarco, and S. Kamin (2011), “International Capital Flows and the Returns to Safe Assets in the United States, 2003-2007,” Board of Governors of the Federal Reserve System International Finance Discussion Paper 1014.
- Boppart, T. (2014) “Structural Change and the Kaldor Facts in a Growth Model with Relative Price Effects and Non-Gorman Preferences,” *Econometrica*, 82, 2167–2196.
- Buera, F., and J. Kaboski (2009), “Can Traditional Theories of Structural Change Fit the Data?” *Journal of the European Economic Association*, 7, 467–477.
- Caballero, R., E. Farhi, and P.-O. Gourinchas (2008), “An Equilibrium Model of ‘Global Imbalances’ and Low Interest Rates,” *American Economic Review*, 98, 358–393.
- Chinn, M. D., and H. Ito (2007), “Current Account Balances, Financial Development and Institutions: Assaying the World ‘Saving Glut’,” *Journal of International Money and Finance*, 26, 546–569.

- Congressional Budget Office (2012), “The 2012 Long-Term Budget Outlook,” www.cbo.gov/publication/43288.
- Congressional Budget Office (2015), “The 2015 Long-Term Budget Outlook,” www.cbo.gov/publication/50250.
- Devereux, M. P., R. Griffith, and A. Klemm (2002), “Corporate Income Tax Reforms and International Tax Competition,” *Economic Policy*, 35, 451–495.
- Echevarria, C. (1995), “Agricultural Development vs. Industrialization: Effects of Trade,” *Canadian Journal of Economics*, 28, 631–647.
- Fernandez de Cordoba G., and T. J. Kehoe (1999), “Capital Flows and Real Exchange Rate Fluctuations Following Spain’s Entry into the European Community,” unpublished manuscript.
- Gourinchas, P.-O., and H. Rey (2007), “From World Banker to World Venture Capitalist: U.S. External Adjustment and the Exorbitant Privilege,” in R. H. Clarida, ed., *G7 Current Account Imbalances: Sustainability and Adjustment*, University of Chicago Press, 11–66.
- Greenspan, A. (2005), “Monetary Policy Report to the Congress,” statement before the Committee on Banking, Housing and Urban Affairs, U.S. Senate, February 16.
- Gruber, J. W., and S. B. Kamin (2007), “Explaining the Global Pattern of Current Account Imbalances,” *Journal of International Money and Finance*, 26, 500–522.
- Herrendorf, B., R. Rogerson, and A. Valentinyi (2013), “Two Perspectives on Preferences and Structural Transformation,” *American Economic Review*, 103, 2752–2789.
- Jin, K. (2012), “Industrial Structure and Capital Flows,” *American Economic Review*, 102, 2111–2146.
- Kehoe, T. J., and K. J. Ruhl (2009), “Sudden Stops, Sectoral Reallocations, and the Real Exchange Rate,” *Journal of Development Economics*, 89, 235–249.
- King, M. and D. Low (2014), “Measuring the ‘World’ Real Interest Rate,” NBER Working Paper 19887.
- Kongsamut, P., S. Rebelo, and D. Xie (2001), “Beyond Balanced Growth,” *Review of Economic Studies*, 68, 869–882.
- Krishnamurthy, A., and A. Vissing-Jorgensen (2007), “The Demand for Treasury Debt,” NBER Working Paper 12881.
- Lane, P. R., and G. M. Milesi-Ferretti (2007), “The External Wealth of Nations Mark II: Revised and Extended Estimates of Foreign Assets and Liabilities, 1970–2004,” *Journal of International Economics*, 73, 223–250.

- Lane, P. R., and G. M. Milesi-Ferretti (2008), “Where Did All the Borrowing Go? A Forensic Analysis of the U.S. External Position,” IMF Working Paper WP/08/28.
- McGrattan, E.R. and E.C. Prescott (2000), “Is the Stock Market Overvalued?” *Federal Reserve Bank of Minneapolis Quarterly Review*, 24, 20–40.
- McGrattan, E.R. and E.C. Prescott (2005), “Taxes Regulations, and the Value of U.S. and U.K. Corporations,” *Review of Economic Studies*, 72, 767–796.
- Matsuyama, K. (1992), “Agricultural Productivity, Comparative Advantage, and Economic Growth,” *Journal of Economic Theory*, 58, 317–334.
- Matsuyama, K. (2009), “Structural Change in an Interdependent World: A Global View of Manufacturing Decline,” *Journal of the European Economic Association*, 7, 478–486.
- Mendoza, E. G., V. Quadrini, and J.-V. Ríos-Rull (2009), “Financial Integration, Financial Development, and Global Imbalances,” *Journal of Political Economy*, 117, 371–416.
- Ngai, L. R., and C.A. Pissarides (2007), “Structural Change in a Multisector Model of Growth,” *American Economic Review*, 97, 429–443.
- Obstfeld, M., and K. Rogoff (2009), “Global Imbalances and the Financial Crisis: Products of Common Causes,” CEPR Discussion Paper 7606.
- Perri, F., and A. Fogli (2010), “Macroeconomic Volatility and External Imbalances,” unpublished manuscript.
- Scott, R. E. (2015), “Manufacturing Job Loss: Trade, Not Productivity, Is the Culprit,” EPI Issue Brief 402.
- Sposi, M. (2012), “Evolving Comparative Advantage, Structural Change, and the Composition of Trade,” unpublished manuscript.
- Timmer, M. P., Dietzenbacher, E., Los, B., Stehrer, R., and de Vries, G. J. (2015), “An Illustrated User Guide to the World Input-Output Database: The Case of Global Automotive Production,” *Review of International Economics*, 23, 575–605.
- United Nations, Department of Economic and Social Affairs, Population Division (2011), “World Population Prospects: The 2010 Revision, Volume I: Comprehensive Tables,” ST/ESA/SER.A/313.
- Uy, T., K.-M. Yi, and J. Zhang (2013), “Structural Change in an Open Economy,” *Journal of Monetary Economics*, 60, 667–682.
- Warnock, F. E., and V. C. Warnock (2009), “International Capital Flows and U.S. Interest Rates,” *Journal of International Money and Finance*, 28, 903–919.

Table 1: 1992 Input-output matrix (U.S. GDP = 100)

		Intermediate inputs				Final demand				Gross output	
		USA			ROW		USA			ROW	
		Goods	Services	Construction	Goods	Services	Consumption	Government	Investment	Consumption	
USA	Goods	21.52	9.96	3.14	3.10	1.14	7.66	1.91	5.26	2.72	56.41
	Services	11.74	39.23	2.99	0.90	1.35	54.97	13.68	3.01	0.47	128.33
	Construction	0.00	0.00	0.00	0.00	0.00	0.00	0.00	10.05	0.00	10.05
ROW	Goods	3.17	1.11	0.33	89.56	41.13	1.64	0.41	1.68	82.31	221.33
	Services	0.44	1.08	0.10	42.59	88.67	0.19	0.05	0.02	205.37	338.51
Value added		19.55	76.96	3.50	85.18	206.22	0.00	0.00	0.00	0.00	391.40
Gross output		56.41	128.33	10.05	221.33	338.51	64.47	16.05	20.02	290.87	

Table 2: Calibration

Parameter	Value	Target
Gross output parameters		
$1/(1-\eta), 1/(1-\xi)$	0.05, 0.03	Atalay (2014)
$1/(1-\zeta_i)$	3.0, 1.0	International macro literature
$\Lambda_i^l, A_i^l, \Pi_i^l$	data appendix	1992 gross output, value added, intermediates
$\lambda_i^{us}, \lambda_i^{rw}$	data appendix	1992 value added shares in gross output
$\pi_{ij}^{us}, \pi_{ij}^{rw}$	data appendix	1992 shares of goods and services in intermediate use
$\mu_{ij}^{us}, \mu_{ij}^{rw}$	data appendix	1992 country shares in intermediate use
α_i^{us}	0.33, 0.35, 0.17	Aggregate capital share of 0.34, 1992 sectoral labor shares
$\{\bar{Y}_{it}^{us}, \bar{Y}_{it}^{rw}\}_{t=1992}^{2092}$	data appendix	Value added and labor compensation (BEA)
Final demand parameters		
$1/(1-\rho), 1/(1-\nu)$	0.65, 0.65	Atalay (2014)
$1/(1-\nu)$	1.00	Bems (2008)
$1/(1-\sigma_i)$	2.00, 1.00	International macro literature
$\Theta_i^{ush}, \Theta_i^{usg}, \Theta_i^{usx}, \Theta_i^{rw}$	data appendix	1992 final demand levels
$\varepsilon_i^{ush}, \varepsilon_i^{usg}, \varepsilon_i^{usx}, \varepsilon_i^{rw}$	data appendix	1992 sectoral shares of final demand
$\theta_i^{ush}, \theta_i^{usg}, \theta_i^{usx}, \theta_i^{rw}$	data appendix	1992 country shares of final demand by sector
Household and government parameters		
β^{us}, β^{rw}	0.99, 0.99	Long-run interest rate, capital/GDP
$1/(1-\psi)$	0.50	Intertemporal elasticity of substitution = 0.5
ϕ^{us}, ϕ^{rw}	0.33, 0.33	1992 leisure time = 2/3
$\bar{c}_g^{us}, \bar{c}_s^{us}$	1.87, 13.79	Herrendorf, Rogerson, and Valentinyi (2013)
$\bar{c}_g^{rw}, \bar{c}_s^{rw}$	8.25, 60.79	Herrendorf, Rogerson, and Valentinyi (2013)
$\{n_t^{us}, \bar{\ell}_t^{us}, n_t^{rw}, \bar{\ell}_t^{rw}\}_{t=1992}^{2092}$	data appendix	World Population Prospects: 2010 Revision
$\{\bar{c}_t^{usg}, \bar{b}_t^{usg}\}_{t=1992}^{2092}$	data appendix	BEA and CBO
Capital formation parameters		
k_{1992}^{us}	277.9	Backus, Henriksen, and Storesletten (2007)
δ	0.05	Fixed capital consumption, 1951–2004 (BEA)
τ_k^{us}	0.39	Devereux, Griffith, and Klemm (2002)

Note: The data appendix can be found in the online supplementary materials.

Table 3: Model and data, non-targeted moments

Statistic	Data	Model
Intermediate share of U.S. gross output, 1995	44.71	48.24
Intermediate share of U.S. gross output, 2011	43.68	46.75
<i>Change</i>	-1.04	-1.50
Goods share of U.S. intermediates, 1995	43.21	40.03
Goods share of U.S. intermediates, 2011	36.40	34.63
<i>Change</i>	-6.81	-5.40
Goods share of U.S. consumption, 1995	15.30	13.73
Goods share of U.S. consumption, 2011	13.63	11.04
<i>Change</i>	-1.66	-2.68
Goods share of U.S. investment, 1995	40.46	34.66
Goods share of U.S. investment, 2011	40.26	34.66
<i>Change</i>	-0.19	0.00
Std. dev. of U.S. goods trade balance, 1992–2012	1.54	1.45
Std. dev. of U.S. services trade balance, 1992–2012	0.24	0.20

Table 4: Goods sector's employment share decline in baseline model vs. alternatives

Model	Benchmark	No-saving-glut counterfactual	Difference
Baseline	99.16	84.06	15.10
<i>Sensitivity analyses</i>			
Trade wedges	94.63	84.06	10.57
Homothetic preferences	93.53	79.10	14.42
No intermediates	66.16	53.40	12.75
Same GO/VA in all US sectors	86.14	75.00	11.14

Note: GO/VA is the gross output to value added ratio.

Figure 1: U.S. trade balance vs. goods sector's employment share




Figure 2a: Labor productivity in goods, services, and construction




Figure 3: U.S. trade balance, current account balance, and real exchange rate




Figure 4: U.S. trade balance in baseline model vs. data

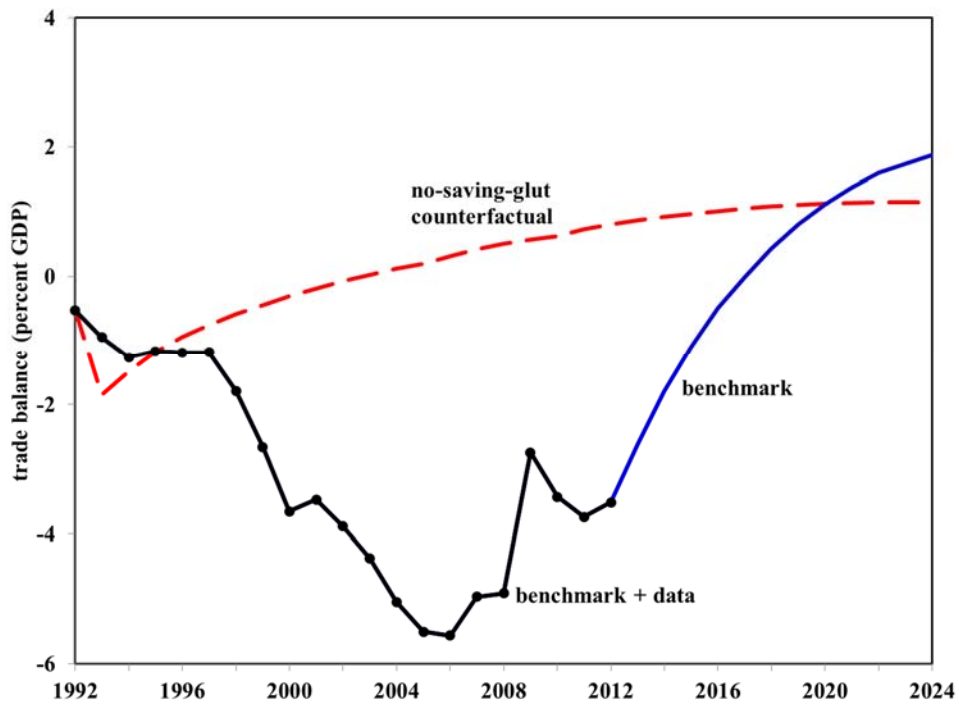


Figure 5: U.S. real exchange rate in baseline model vs. data




Figure 6: Disaggregated trade balances in baseline model vs. data




Figure 7: Goods and construction employment shares in baseline model vs. data

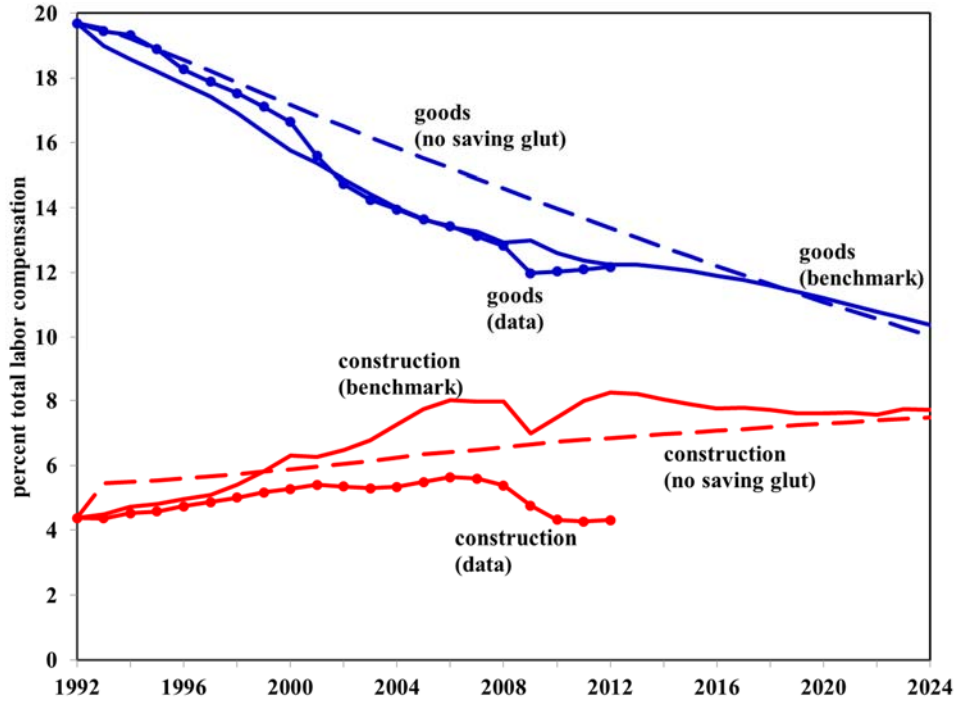


Figure 8: Goods employment share in trade wedge model vs. baseline model




Figure 9: Investment in domestic-saving-drought model vs. baseline model




Figure 10: U.S. real exchange rates with China and other trade partners




Figure 11: U.S. real interest rate in baseline model vs. data

