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 Minnesota
and Federal Reserve Bank of Minneapolis

May 2013

ABSTRACT

The United States has borrowed heavily from the rest of the world since the early 1990s. During this period, the share of employment in goods production has fallen dramatically. We build a general equilibrium model of the United States and the rest of the world to assess the quantitative impact of U.S. borrowing on goods-sector employment. We find that U.S. borrowing is not the most important driver of the decline in goods-sector employment, instead it is the rapid productivity growth in goods production. As the U.S. repays its debt to the rest of the world, its trade balance will reverse, but goods-sector employment will continue to fall. A sudden stop in foreign lending in 2015–2016 would cause a sharp trade balance reversal and painful reallocation across sectors, but would not affect long-term structural change.

Keywords: global imbalances, real exchange rate, structural change
JEL Classification: E13, F34, O41

---

* We thank David Backus, Kei-Mu Yi, and Frank Warnock for helpful discussions. Seminar participants at PUC-Rio, University of Pittsburgh, and University of California, Davis, as well as conference participants at ITAM and at Bogazici University made useful comments 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 presented in the figures are available at http://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, households and the government in the United States borrowed heavily from the rest of the world. As U.S. borrowing grew, the U.S. net international investment position deteriorated by 4.0 trillion dollars, and by 2012, the United States owed the rest of the world 4.4 trillion dollars. In this paper, we use a dynamic, general equilibrium model of the United States and the rest of the world to study what will happen when foreigners' willingness to lend wanes and the United States begins to repay its debt. Our model is capable of addressing a large number of questions about this repayment process, but we focus on two questions that have particular relevance for policymakers. First, will employment return to goods-producing sectors when U.S. borrowing ends and trade deficits become trade surpluses? Second, how would a disorderly end to borrowing — a sudden stop — affect goods-sector employment dynamics?

To answer these questions, we must take a stand on the source of U.S. borrowing. One of the most common explanations is that foreign demand for saving increased during the past two decades, making them more willing to trade their goods and services for U.S. bonds. Ben Bernanke (2005), the Chairman of the Board of Governors of the U.S. Federal Reserve, coined the term *global savings glut* to refer to this idea. In this paper we adopt Bernanke's savings glut hypothesis. The literature proposes a variety of explanations for increased demand for saving in the rest of the world, particularly in developing countries like China, such as a lack of financial development (Caballero, Farhi, and Gourinchas, 2008; Mendoza, Quadrini, and Rios-Rull, 2009), differences in business cycle or structural growth properties (Backus, Henriksen, Lamber, and Telmer, 2006; Perri and Fogli, 2010), and demographic differences (Du and Wei, 2010). We do not take a stand on which of these explanations, if any, are correct. Instead, we take the savings glut as given and study its impact on the U.S. economy by calibrating a process for the preferences of households 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 between 1992 and 2012. Our model closely captures several other key facts about the U.S. economy since the early 1990s. We view our model's performance along this dimension as evidence in favor of the savings glut hypothesis and an indication that our model is well-suited to provide quantitatively accurate answers to the questions posed above.

There is a commonly-held belief in policy circles that U.S. trade deficits have played an important role in the decline of employment in the U.S. goods sector and that an end to these deficits will reverse a large part of this trend (see, for example, Bivens, 2006; Scott, Jorgensen, and Hall, 2013). The intuition is simple: Imported goods are substitutes for those produced in the domestic goods sector. As the United States trades bonds for foreign goods, the shift away from domestically produced goods frees up labor from that sector, which is reallocated to producing services and construction. When the debt has to be repaid, labor will flow back into the goods sector to produce the extra goods needed to repay the debt. Contrary to this belief, we show that this mechanism has not been the major determinant in the decline in goods-sector employment, and that this decline is likely to continue even after the savings glut ends.

It is easy to see why some believe that U.S. goods-sector employment is driven by foreign lending. Figure 1 shows that the share of employment in the goods-sector — agriculture, mining, and manufacturing — has fallen dramatically just as the trade deficit has grown. The correlation between the level of the trade balance and our preferred measure of the goods employment share during this period is more than 0.8, but as we will show in this paper, only a small fraction of the decline in the goods-sector employment share is attributable to the savings glut. The bulk of this decline is accounted for by a standard force in the structural change literature: Labor productivity has consistently grown faster in the goods sector than in other sectors, as can be seen in figure 2. The importance of differential productivity growth means that the end of the savings glut will not bring about a resurgence of goods-sector employment; as long as the productivity trends in figure 2 persist, the goods-sector employment share will continue to decline.

Our research is related to the structural change literature, which has long emphasized asymmetric productivity growth rates as an important driver of long-run reallocation of labor across sectors. Several recent studies embed this mechanism, originally due to Baumol (1967), into growth models that are consistent with aggregate balanced growth (Ngai and Pissarides, 2007; Buera and Kaboski, 2009). We take a similar approach, using data on value added and labor compensation in the goods, services and construction sectors in the United States over 1987–2011 to calibrate the productivity growth rates in each of these sectors. Labor productivity in the goods sector grew at a substantially higher rate over this period than in the other two sectors, and this plays a large role in causing employment in the goods sector to decline in our model. Several recent papers incorporate structural change into open-economy models to study the importance of trade for long-run compositional changes (Echevarria, 1995; Matsuyama, 1992 and 2009; Ui, Yi, and Zhang, 2012; Sposi, 2012). With the exception of Sposi (2012), these studies use models of balanced trade, abstracting from international capital flows. We place capital flows at the forefront of our analysis, and we perform a quantitative assessment of the relative contributions of traditional structural change forces (asymmetric labor productivity growth) and the savings glut to the decline in goods-sector employment in the United States.

In our baseline model, we assume that the transition from current account deficit to surplus will happen gradually. In an extension of the model, we consider the effects of a sharp reversal of the current account. The possibility of a U.S. "sudden stop" arises frequently in policy discussions. Bernanke, for example, argued in a 2005 speech to the Virginia Association of Economists, that while a gradual rebalancing process is likely, he cannot rule out a sudden stop event: "[T]he underlying sources of the U.S. current account deficit appear to be medium-term or even long-term in nature, suggesting that the situation will eventually begin to improve, although a return to approximate balance may take some time. Fundamentally, I see no reason why the whole process should not proceed smoothly. However, the risk of a disorderly adjustment in financial markets always exists." Sudden stops in the past have often been accompanied by severe economic disruption in the short term: abrupt current account and trade balance reversals, large real exchange rate depreciations, substantial productivity-driven output contractions, and sharp reallocations of factors across sectors (Calvo, Izquierda, and Mejía, 2004; Calvo, Izquierda, and Talvi, 2005; Kehoe and Ruhl, 2009). We use our model to study what a sudden stop in the United States in 2015–2016 might look like and quantify the cost of the associated disruption in economic activity. Our results indicate that in many respects, a sudden stop in the United States would look similar to historical episodes in emerging economies. The trade balance would rise by more than 4 percent on impact, the real exchange rate would depreciate by 13 percent, and employment would shift from away from nontraded sectors (construction) to traded sectors (goods and services). Also as in sudden stops of the past, these effects would be short-lived; several years after the sudden stop ends the U.S. economy would be in almost exactly the same state it would have been if the sudden stop had never happened at all.

Our results suggest that the impact of a sudden stop will be temporary, but that this does not mean the welfare costs of a sudden stop are small. We construct a measure of the real income of U.S. households in 1992 to compare the welfare of U.S. households in three scenarios: the baseline in which the savings glut ends in gradual rebalancing, an alternative in which it ends in a sudden stop, and a counterfactual in which the savings glut never happens at all. The savings glut itself markedly improves welfare; life-time real income of U.S. households in 1992 would have been almost $700 billion lower (more than 10 percent of 1992 GDP) if the savings glut had not happened. A sudden stop in 2015–2016, though, would reduce real income from the perspective of model agents in 1992 by more than $1 trillion, completely wiping out the welfare gains created by the savings glut. In other words, if the savings glut ends in a sudden stop, U.S. households would prefer that the savings glut had never happened at all.

Addressing questions about the sectoral allocation of resources requires a good model of the input-output structure of the economy. We include several features in the model that make it uniquely well-suited to address these issues in a rigorous, quantitative fashion. First, we model an economy with three sectors: goods, services, and construction. Both goods and services can be traded, which allows us to capture the fact that the U.S. consistently runs a substantial trade surplus in services. Most multi-sector models in international macroeconomics treat goods — manufacturing, mining, and agriculture — as the only traded sector, and lump all other sectors into a single non-tradable sector. This assumption is at odds with the data: Services are an important component of U.S. exports. Second, we allow the elasticity of substitution between foreign and domestic inputs to differ across sectors. Our calibration assigns a higher elasticity of substitution between domestic and foreign inputs in the goods sector to match the fact that the goods trade balance is significantly more volatile than the services trade balance (see figure 4). Construction is the only non-tradable sector in our model and is used almost entirely to produce investment goods, which means that construction is more sensitive than other sectors to the effects of capital flows and economic fluctuations in general. Third, and perhaps most importantly, we allow labor productivity to grow at different rates across sectors, which allows us to match the fact that labor productivity in the goods sector grew at a faster pace than in other sectors over the past two decades.

Our model highlights two puzzles related to the price effects of the savings glut. First, why did U.S. borrowing continue to increase once the U.S. real exchange rate began to depreciate? In other words, why did U.S. purchases of foreign goods and services continue to increase once foreign goods stopped getting cheaper and started getting more expensive? Figure

3 illustrates this probelm with the timing. After appreciating by more than 20 percent, the U.S. real exchange rate began to depreciate in 2002, but the U.S. trade deficit continued to rise until 2006. While out model matches the magnitude of the real exchange rate appreciation between 1992 and 2011, it fails to generate this timing — the real exchange rate depreciates more slowly in the model than the data. The second puzzle concerns U.S. real interest rates. Bernanke (2005), among others, has argued that demand for saving in the rest of the world is an important factor in explaining persistently low U.S. real interest rates. In our model, foreign demand for saving has very little impact on U.S. real interest rates. The intuition for this result is straightforward: U.S. borrowing primarily financed purchases of foreign goods, and these foreign goods represent a small portion of the U.S. consumption basket. Our result is consistent with the empirical estimates on the impact of foreign purchases of U.S. assets on U.S. real interest rates in Warnock and Warnock (2008) and Krishamurthy and Vissing-Jorgensen (2008). This suggests we may need to look to other factors, like the domestic developments in housing and financial markets discussed by Obstfeld and Rogoff (2009) and Bernanke, Bertaut, DeMarco, and Kamin (2011), to explain low U.S. real interest rates.

We have performed extensive sensitivity analysis of our model. Our main results are robust to a wide range of modeling changes, like altering our assumptions about the path of government spending and abandoning perfect foresight in favor of a stochastic model with rational expectations. The input-output production structure in the U.S. economy and the tradability of services, however, play important quantitative roles in our model's predictions for goods-sector employment. Removing the input-output structure raises the elasticity of substitution between goods and services in gross output, which causes our to model capture a significantly smaller fraction of the decline in goods-sector employment we observe in the data during 1992–2011. This change also leads our model to overstate the role of the savings glut in reducing goods-sector employment during this period; making services non-tradable has a similar effect. To justify our assumption that the U.S. trade deficit has been driven by the rest of the world's demand for saving, we also study a version of our model in which we calibrate the preferences of U.S. households, rather than those in the rest of the world, to match the U.S. trade balance. While this change does not significantly affect our main results about goods-sector employment and the impact of a sudden stop, it leads to predictions for investment and construction employment that are sharply at odds with the data. We believe that these results vindicate our approach to modeling the source of U.S. trade deficits and lend credence to Bernanke's savings glut hypothesis.

The rest of the paper proceeds as follows. Section 2 documents several key facts about the U.S. economy since 1992. Section 3 describes our model environment. Section 4 provides an outline of our quantitative strategy and section 5 details our calibration and data sources for exogenous processes like technological and demographic change. Section 6 presents our main results. We conduct a sensitivity analysis in section 7. Section 8 provides concluding remarks.

2. The U.S. economy during the 1992–2012 savings glut period

In this section we highlight what we view as the key facts about the U.S. economy during the past two decades of sustained borrowing from the rest of the world. Figure 3 illustrates just how much borrowing the United States has done. While the current account balance measures the exact magnitude of capital flows into the United States during this period, we see that the trade balance tracks the current account balance almost exactly. Our model is not equipped to accurately capture the difference between these two series and the trade balance has an exact model analogue, so 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, eventually reaching 5.75 percent of GDP, after which it began to shrink as the financial crisis of 2008–2009 approached. In 2012 the trade deficit remained at more than 3.5 percent of GDP. We view the path of the trade balance as "fact zero" about the savings glut period — this time series is what defines the savings glut in our model.

Fact 1: The real exchange rate appreciated, then depreciated

Figure 3 also presents our first key fact. The figure shows that the U.S. real exchange rate was quite volatile and tracked the trade balance quite closely during the savings glut period. 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 do not significantly affect our results). This approach forms the basis of our conception of the "rest of the world" in our model. The real exchange rate appreciated by almost 22 percent between 1992 and 2002, after which it began to depreciate, ending up almost 32 percent higher than its low in 2002. 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. In our model, this logic holds exactly — the trade balance and real exchange rate move up and down in concert. Given this logic, the lag between the trade balance and the real exchange rate that we see in the figure is puzzling — why did U.S. households continue to increase their purchases of foreign goods and services between 2002 and 2006 once they started to become more expensive? This phenomenon, often referred to as the "J-curve" (see, for example, Backus, Kehoe, and Kydland, 1994), has been extensively documented in the literature. One common explanation is that households and firms cannot adjust their purchases of foreign goods and services immediately, so the trade balance takes time to adjust to relative price movements. Our view is that this kind of story is insufficient to explain the particularly long lag observed in the figure (there are 4 years between the troughs of the two series), but addressing this feature of the data is outside the scope of our current paper.

Fact 2: The goods sector drove aggregate the trade balance dynamics

Figure 4 presents our second key fact. Here we disaggregate the trade balance, plotting the trade balances for goods and services (as defined in the national income and product accounts) separately. We see that the goods trade balance tracks the aggregate trade balance closely, while the services trade balance remains relatively constant between 0.5 and 1.25 percent of GDP. The fact that the United States has consistently run a trade surplus in services provides the basis for one of the key contributions of our modeling framework. Standard modeling conventions in international macroeconomics lump services into a nontraded 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, and we calibrate home bias parameters for each sector separately to capture the fact that the United States consistently runs a surplus in services and a deficit in goods. As we will see, this fact has important implications for how the end of the savings glut will play out; our model predicts that the services sector will generate the entire trade surplus needed for the United States to repay its debt.

Figure 4 contains one more piece of information that informs our modeling choices; namely, the fact that the services trade surplus is less volatile than the goods trade surplus. As we describe below, we calibrate a lower elasticity of substitution between domestic and foreign inputs in services sector to capture the fact that even as the U.S. real exchange rate fluctuated substantially during the period under observation, the share of domestic inputs in U.S. services production did not move much. We set the elasticities in goods and services to 3 and 1, respectively, so that the average elasticity (weighted by sector size) remains within the standard range of 1.5 and 2.

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

Figure 5 presents our third fact, which we have already highlighted in the previous section. Between 1992 and 2011 the fraction of total labor compensation paid to the goods sector fell from 19.7 percent to 12.4 percent. This is our preferred measure of the goods sector's employment share because it maps directly into our model. As the figure shows, this measure moves closely with more common measures like the share of employees in the goods sector. The construction sector's share of labor compensation rose from 4.4 percent in 1992 to 5.7 percent in 2006, representing the boom in construction employment prior to the financial crisis in 2008–2009. Employment in construction then started to fall once the crisis hit, and by 2012 the construction sector's share of labor compensation was at 4.4 percent, returning to its 1992 level. We refer to this trend as the "boom" and "bust" in construction employment.

Trends in labor productivity across sectors

In addition to the three facts described above, we wish to highlight one additional trend that plays an important role in our results: The rates of labor productivity growth in the goods, services, and construction sectors. We do not refer to this as one of our "key facts," however, since we will not compare our model's results in this dimension with the data; we treat labor productivity growth as an exogenous parameter in our model. Figure 2 plots labor productivity in each of these three sectors since 1987. The average growth rate of labor productivity in the goods sector during this period is 4.3 percent. By comparison, the services and construction sectors saw average labor productivity growth rates of 1.3 percent and -0.4 percent, respectively. We use these average growth rates in our model. The fact that labor productivity in goods grew so much more rapidly than in the rest of the economy is the main driver of the decline in goods sector employment in our model. This is a standard force in the structural change literature and has an obvious connection with employment — compared to other sectors, U.S. goods producers needed fewer and fewer workers to supply demand from U.S. households and firms.

3. Model

We model an economy with two countries, the United States and the rest of the world (RW). Throughout this section, we use the superscripts *us* and *rw* to denote prices, quantities, and parameters in the United States and the rest of the world. The length of a period is one year. Each country has a representative household that works, consumes, and saves to maximize utility subject to a sequence of budget constraints. We assume that the only internationally traded assets are one-period bonds denominated in units of the U.S. consumer price index (CPI). Households have perfect foresight over the future trajectory of the world economy — there is no uncertainty. Each country produces several commodities that serve both intermediate and final uses. We model the U.S. production structure in detail, using an input-output structure which we calibrate to a benchmark input-output matrix published by the U.S. Bureau of Economic Analysis (BEA). We model production in the rest of the world in a simpler fashion, abstracting from investment and domestic input-output linkages. We model the U.S. government in a somewhat reduced-form fashion as well. The government's spending and debt as fractions of U.S. GDP are specified exogenously, and the government levies lump-sum taxes on U.S. households to ensure its budget constraints are satisfied.

Production

The United States produces 4 commodities: goods y_{gt}^{us}, services, y_{st}^{us}, construction y_{ct}^{us}, and investment y_{it}^{us}. The rest of the world produces its own goods, y_{gt}^{rw}, and services, y_{st}^{rw}. The subscripts *g*, *s*, *c*, and *i* denote sectors: goods, services, construction, and investment. The price of each commodity uses similar notation; for example, the price of U.S. investment is p_{it}^{us}. All commodities are sold in perfectly competitive markets, and sectors in the United States are linked together through an input-output structure. In what follows, we provide a detailed description of production of each type of commodity, starting with U.S. production then moving on to the rest of the world.

In the U.S. economy, each commodity $j = g, s$ is produced using capital k_{jt}^{us} and labor ℓ_{jt}^{us}, along with intermediate inputs of U.S. goods z_{gjt}^{us}, U.S. services z_{sjt}^{us}, U.S. construction z_{cjt}^{us}, and imported intermediate inputs m_{jt}^{us} purchased from the same sector j in the rest of the world. The production functions take the form

$$(1)\quad y_{jt}^{us} = M_j^{us}\left\{\mu_j^{us}\left(\min\left[\frac{z_{gjt}^{us}}{a_{gj}^{us}},\frac{z_{sjt}^{us}}{a_{sj}^{us}},\frac{z_{cjt}^{us}}{a_{cj}^{us}},A_j^{us}\left(k_{jt}^{us}\right)^{\alpha_j}\left(\gamma_{jt}^{us}\ell_{jt}^{us}\right)^{1-\alpha_j}\right]\right)^{\zeta_j} + (1-\mu_j^{us})\left(m_{jt}^{us}\right)^{\zeta_j}\right\}^{\frac{1}{\zeta_j}}.$$

This nested production function embeds a Leontief input-output structure in a standard Armington aggregator. The parameters of the production functions are as follows: M_j^{us} and A_j^{us} are constant scaling factors used to facilitate calibration. μ_j^{us} governs the share of imports in production. ζ_j governs the elasticity of substitution between domestic and imported inputs. $\left(a_{gj}^{us}, a_{sj}^{us}, a_{cj}^{us}\right)$ govern the shares of goods, services, and construction in gross output. γ_{jt}^{us} is sector-specific labor productivity. α_j is capital's share in value added.

We allow the Armington elasticities ζ_j to differ across sectors in order to capture the fact that the goods trade balance is considerably more volatile than the services trade balance (see figure 3). We also allow labor productivity γ_{jt}^{us} to grow at different rates across sectors to capture the fact that productivity in the goods sector has grown faster than in other sectors. We discuss these issues in more detail in the calibration section.

Consistent with our input-output table, the construction sector is the only purely nontraded sector. The production function in the construction sector is a special case of (1) where the share of imported inputs in production is set to zero:

$$(2)\qquad y_{ct}^{us} = \min\left[\frac{z_{gct}^{us}}{a_{gc}^{us}},\frac{z_{sct}^{us}}{a_{sc}^{us}},\frac{z_{cct}^{us}}{a_{cc}^{us}},A_c^{us}\left(k_{ct}^{us}\right)^{\alpha_c}\left(\gamma_{ct}^{us}\ell_{ct}^{us}\right)^{1-\alpha_c}\right].$$

U.S. producers in all three sectors $j \in \{g,s,c\}$ choose inputs of intermediates and factors to minimize costs, which implies standard marginal product pricing conditions for capital and labor.

The U.S. investment good is produced using inputs z_{git}^{us}, z_{sit}^{us}, and z_{cit}^{us} of composite goods, composite services, and construction (all domestic) according to a Cobb-Douglas technology:

$$(3)\qquad y_{it}^{us} = G\left(z_{git}^{us}\right)^{\theta_g}\left(z_{sit}^{us}\right)^{\theta_s}\left(z_{cit}^{us}\right)^{\theta_c},\quad \theta_g+\theta_s+\theta_c = 1.$$

Our Cobb-Douglas specification is consistent with empirical evidence reported by Bems (2008), who shows that expenditure shares on investment inputs are approximately constant over time across a range of countries.

We model the rest of the world's production structure in less detail, abstracting from investment and input-output linkages. Goods and services in the rest of the world are produced using labor ℓ_{jt}^{rw} and imported intermediate inputs m_{jt}^{rw} from the same sector j in the United States. The production functions are simpler nested Armington aggregators of the form

$$(4) \qquad y_{jt}^{rw} = M_j^{rw}\left(\mu_j^{rw}\left(\gamma_{jt}^{rw}\ell_{jt}^{rw}\right)^{\zeta_j} + (1-\mu_j^{rw})\left(m_{jt}^{rw}\right)^{\zeta_j}\right)^{\frac{1}{\zeta_j}}.$$

The parameters have the same interpretation as described above.

Households

Each country is populated by a continuum of identical households. We draw a distinction between the total and working-age populations as these two groups grow at different rates. We denote the total U.S. population by $\tilde{n}_t^{us}$ and the working-age population by $\overline{\ell}_t^{us}$. We evaluate consumption per capita on an adult-equivalent basis, defining the U.S. adult-equivalent population as

$$(5) \qquad n_t^{us} = \overline{\ell}_t^{us} + (\tilde{n}_t^{us} - \overline{\ell}_t^{us})/2.$$

The rest of the world's demographic variables are defined similarly. As in the description of the economy's production structure, we begin with the problem of a U.S. household then move on to the rest of the world.

We normalize the amount of time available for work and leisure by a U.S. working-age person to one and denote total U.S. labor supply by ℓ_t^{us}. U.S. households choose labor supply, consumption of composite goods and services, c_{gt}^{ush} and c_{st}^{ush}, investment i_t^{us}, and bond holdings b_t^{us}, to maximize utility

$$(6) \qquad \sum_{t=0}^{\infty}\beta^t\left(\left(\varepsilon^{ush}\left(\frac{c_{gt}^{ush}}{n_t^{us}}\right)^{\rho} + (1-\varepsilon^{ush})\left(\frac{c_{st}^{ush}}{n_t^{us}}\right)^{\rho}\right)^{\frac{\eta\psi}{\rho}}\left(\frac{\overline{\ell}_t^{us} - \ell_t^{us}}{\overline{\ell}_t^{us}}\right)^{(1-\eta)\psi} - 1\right)/\psi,$$

subject to the budget constraints

(7) $$p_{gt}^{us}c_{gt}^{ush} + p_{st}^{us}c_{st}^{ush} + p_{it}^{us}i_t^{us} + q_t b_{t+1}^{ush} = w_t^{us}\ell_t^{us} + p^{us}(p_{gt}^{us}, p_{st}^{us})b_t^{ush} + (1-\tau_k^{us})r_{kt}^{us}k_t^{us} - T_t^{us},$$

the law of motion for capital

(8) $$k_{t+1}^{us} = (1-\delta)k_t^{us} + i_t^{us},$$

appropriate non-negativity constraints, initial conditions for the capital stock and bond holdings $\bar{k}_0^{us}$ and $\bar{b}_0^{us}$, and a constraint on bond holdings that rules out Ponzi schemes but does not otherwise bond in equilibrium. We use the superscript ush for U.S. households' consumption and bond holdings to distinguish them from those of the U.S. government, for which we use the superscript usg. Notice that we also g to distinguish goods from services and construction.

Bonds are denominated in units of the U.S. CPI, defined as

(9) $$p^{us}(p_{gt}^{us}, p_{st}^{us}) = \frac{p_{gt}^{us}c_{g1992}^{ush} + p_{st}^{us}c_{s1992}^{ush}}{P_{g1992}^{us}c_{g1992}^{ush} + P_{s1992}^{us}c_{s1992}^{ush}}.$$

We model discount bonds, so the price q_t represents the price in period t of one unit of the U.S. CPI basket in period t+1. The real interest rate in units of the U.S. CPI is given by

(10) $$1 + r_{t+1} = p^{us}(p_{gt}^{us}, p_{st}^{us}) / q_t$$

Households pay constant proportional taxes τ_k^{us} on capital income and a lump-sum tax or transfer T_t^{us}. We use the capital income tax to obtain a sensible calibration for the initial capital stock and depreciation rate. In our calibration we also allow the tax rate on capital income in 1993 to differ from the constant rate in order to match the level of investment in 1992. The first-order conditions for bonds and investment imply a no-arbitrage condition,

(11) $$p^{us}(p_{gt}^{us}, p_{st}^{us}) / q_t = \left((1-\tau^{us})r_{kt+1}^{us} + p_{it+1}^{us}(1-\delta)\right) / p_{it}^{us},$$

which dictates that the return on bonds must equal the return on investing in an additional unit of capital. During the sudden stop episode that may occur in 2015, bond-holdings are fixed and the internal real interest rate is determined endogenously in each country separately.

The rest of the world's households solve a slightly simpler problem. We abstract from investment dynamics in the rest of the world, so the only way the rest of the world's households can save is by buying bonds. Labor supply is still endogenous, however, and the rest of the world's households have similar preferences

(12) $$\sum_{t=0}^{\infty}\beta^t\omega_t^{rw}\left(\left(\varepsilon^{rw}\left(\frac{c_{gt}^{rw}}{n_t^{rw}}\right)^{\rho}+(1-\varepsilon^{rw})\left(\frac{c_{st}^{rw}}{n_t^{rw}}\right)^{\rho}\right)^{\frac{\eta\psi}{\rho}}\left(\frac{\bar{\ell}_t^{rw}-\ell_t^{rw}}{\bar{\ell}_t^{rw}}\right)^{(1-\eta)\psi}-1\right)/\psi\ .$$

The only differences from the U.S. household's utility function are the share parameter ε^{rw} and the parameter ω_t^{rw}. The latter is an intertemporal demand shifter, which we calibrate to match the U.S. trade balance during the savings glut period of 1992–2011. During this period ω_t^{rw} falls, reflecting a reduction in utility gained from consumption during this period compared to future consumption. This is the driving force behind the savings glut. The rest of the world's representative household chooses labor supply ℓ_t^{rw}, consumption of goods and services, c_{gt}^{rw} and c_{st}^{rw}, and bond holdings b_t^{rw} to maximize utility subject to the budget constraints,

(13) $$p_{gt}^{rw}c_{gt}^{rw}+p_{st}^{rw}c_{st}^{rw}+q_t b_{t+1}^{rw}=w_t^{rw}\ell_t^{rw}+p^{us}(p_{gt}^{us},p_{st}^{us})b_t^{rw}\ ,$$

and similar non-negativity and no-Ponzi constraints to those that U.S. households face. The rest of the world's CPI is defined analogous to the U.S. CPI in (9). We then define the real exchange rate as

(14) $$rer_t=p^{rw}(p_{gt}^{rw},p_{st}^{rw})\,/\,p^{us}(p_{gt}^{us},p_{st}^{us})\,.$$

U.S. government

The government in the United States levies taxes and sells bonds to finance exogenously-required expenditures on consumption of composite goods and services. The government's budget constraint is

(15) $$p_{gt}^{us}c_{gt}^{usg}+p_{st}^{us}c_{st}^{usg}+q_t b_{t+1}^{usg}=\tau_k^{us}r_{kt}^{us}K_t^{us}+T_t^{us}+p^{us}(p_{gt}^{us},p_{st}^{us})b_t^{usg}\ ,$$

where, we use the superscript usg to distinguish public from private U.S. consumption and bond holdings. We specify time paths for government consumption expenditures and debt as fractions of GDP, using historical data for 1992—2011 and projections for the future. We allow the lump-sum tax T_t^{us} to vary as necessary to ensure that the government's budget constraint is always satisfied. More formally, let ν_t and υ_t denote the fractions of nominal GDP in period *t* that government consumption expenditures and debt must equal respectively. We require that

(16) $$p_{gt}^{us}c_{gt}^{usg} + p_{st}^{us}c_{st}^{usg} = \nu_t GDP_t^{us}$$

and

(17) $$b_t^{usg} = -\upsilon_t GDP_t^{us} .$$

In this setup we must specify the degree to which the government can substitute goods for services in consumption. We remain neutral as to whether goods and services are complements or substitutes for the government, setting the elasticity of substitution to one. The government therefore chooses c_{gt}^{usg} and c_{st}^{usg} to maximize

(18) $$\left(c_{gt}^{usg}\right)^{\varepsilon^{usg}} \left(c_{st}^{usg}\right)^{1-\varepsilon^{usg}}$$

subject to (16). We assume that government spending does not enter the household's utility function (or equivalently, enters in a separable fashion), nor does it enter any of the production functions.

It is important to point out that our model exhibits near-Ricardian equivalence. Ricardian equivalence breaks down only when we introduce unexpected events — the savings glut and the sudden stop. This implies that the distribution of debt between public and private holdings is not uniquely defined. Unanticipated changes in the time path of government spending and debt that accompany these events do affect the model's equilibrium dynamics, particularly in the short run.

Market clearing and equilibrium

The market clearing conditions for U.S. composite goods and services are

(19) $$c_{jt}^{ush} + c_{jt}^{usg} + m_{jt}^{rw} + z_{gjt}^{us} + z_{sjt}^{us} + z_{cjt}^{us} + z_{jit}^{us} = y_{jt}^{us} ,\ \ j = g, s .$$

Construction is not traded or used for consumption, so the market clearing condition for construction is

(20) $$z_{cgt}^{us} + z_{cst}^{us} + z_{cct}^{us} + z_{cit}^{us} = y_{dt}^{us} .$$

Market clearing for U.S. investment is

(21) $$i_t^{us} = y_{it}^{us} .$$

The market clearing conditions for the rest of the world's composite goods and services are

(22) $$c_{jt}^{rw} + m_{jt}^{us} = y_{tjt}^{rw}, \;\; j = g, s$$

Factor markets must also clear:

(23) $$k_{gt}^{us} + k_{st}^{us} + k_{ct}^{us} = k_{t}^{us}, \quad \ell_{gt}^{us} + \ell_{st}^{us} + \ell_{ct}^{us} = \ell_{t}^{us}, \quad \ell_{gt}^{rw} + \ell_{st}^{rw} = \ell_{t}^{rw}.$$

Finally, the bond market must clear:

(24) $$b_t^{ush} + b_t^{usg} + b_t^{rw} = 0.$$

An equilibrium in our model for a given sequence of time series parameters $\{\omega_t^{rw}, \nu_t, \upsilon_t\}_{t=0}^{\infty}$ and initial conditions $(\bar{b}_0^{ush}, \bar{b}_0^{ush}, \bar{k}_0^{us})$ consists of a sequence of all model variables such that households in the U.S. and the rest of the world maximize their utilities subject to their constraint sets, prices and quantities satisfy marginal product pricing conditions for all 10 commodities, all market clearing conditions are satisfied, and the U.S. government solves its consumption spending allocation problem in each period. When we solve the model numerically, we require that equilibria converge to balanced growth paths after 100 years. There are an infinite number of possible balanced growth paths — one for every combination of public and private bond holdings.

4. Outline of our quantitative strategy

Before we proceed with our calibration and quantitative exercises, we first describe our overall quantitative strategy. Our first step is to calibrate the model's parameters and initial conditions so that the equilibrium in which neither the sudden stop nor savings glut occur replicates the benchmark input-output matrix and national account figures published by the BEA for 1992. We calibrate the rest of the world's preference parameter ω_{1992}^{rw} to match the U.S. trade balance in 1992, and assume that it converges quickly to a constant value of one thereafter. In this step, agents in our model do not expect the savings glut to occur. We do not calibrate our model to match any time series at all in this step. We treat the model's equilibrium dynamics in this scenario as a counterfactual exercise, allowing us to ask the question: what would have happened over the past two decades, and in the future, had the savings glut not happened at all? We could skip this step and go right to the next one without changing our main results significantly, but we view this as the most natural way to think about the agents' expectations in the early 1990s; we think it is extremely unlikely that U.S. households and firms in 1992 foresaw the kind of borrowing the United States would do over the next two decades.

Our second step is to solve for the model's dynamics in the scenario where the savings glut actually happens. Here, we hold fixed all parameters calibrated during the first step, and calibrate the values of ω_t^{rw} for 1993–2011 so that the equilibrium replicates exactly the aggregate U.S. trade balance during this period. As in the first step of our exercise, we assume that ω_t^{rw} gradually converges to one once this period ends. We use the equilibrium values of capital and bond holdings in 1993 from the no-savings glut first step as initial conditions in this second step. The savings glut, which manifests in our model as temporarily reduced utility from consumption and leisure in the rest of the world, is an unanticipated event. Model agents in 1992 do not expect it to occur — its sudden onset is completely unexpected — but they have perfect foresight thereafter. We refer to the model's post-2011 dynamics in this scenario as a "gradual rebalancing," representing the outcome of a slow, orderly end to the forces driving the savings glut. In our sensitivity analysis we relax the perfect foresight assumption and allow for agents to be uncertain about the length of the savings glut.

Our third and final step is to solve for the model's dynamics from 2015 onwards if a sudden stop occurs. We treat a sudden stop as another unforeseen event, and treat the equilibrium state variables from the second step ("gradual rebalancing") in 2015 as the initial conditions. As before, once the sudden stop occurs, model agents have perfect foresight thereafter. We model a sudden stop as a two-year period in which the United States is restricted from borrowing further from the rest of the world, after which foreigners are again willing to purchase U.S. bonds. We assume that the rest of the world's preference parameter ω_t^{rw} converges to its long-run value of one more quickly than in the previous scenario, which we intend to capture the idea that a sudden stop is associated with a faster end to the forces that drove the savings glut in the first place. We also assume that the U.S. government's debt as a fraction of GDP falls to a lower long-run level than in the gradual rebalancing case, representing the idea that a sudden stop is associated with, or perhaps triggers, a long-term change in U.S. government debt policy.

5. Calibration

In the spirit of multisector, static applied general equilibrium models like Kehoe and Kehoe (1994), we calibrate many of the model's parameters so that the equilibrium in 1992 of the model in which the savings glut does not occur replicates the input-output matrix for that same year published by the BEA. There are several discrepancies between the NIPA tables and the input-output matrix, so we let the NIPA tables take precedence and perform several adjustments using the RAS algorithm to the input-output matrix so that it matches the relevant NIPA data exactly. Our adjusted input-output matrix is listed in table 1. The dynamic, open-economy nature of the model introduces several other elements to the model's calibration. We sequence of parameters that govern sector-level productivity growth, demographics, and government consumption and debt are taken directly from the data.

U.S. production parameters

We normalize quantity units so that U.S. GDP is equal to 100 and all prices are equal to one in 1992 — all quantities are expressed as percentages of 1992 GDP. We compute the parameters in the Leontief portion of the U.S. production functions in (1) directly from the input-output matrix. For example, to compute a_{gc}^{us}, the amount of goods needed to produce one unit of gross output in the construction sector, we divide the value in the goods row and construction column (3.79) by gross output in the construction column (10.71). We use a similar procedure to calculate factor shares in value added for each sector. For the Armington aggregators in (1) we first specify values for the elasticities of substitution between domestic and imported inputs. There is some debate over this elasticity as business cycle models tend to imply low elasticities while analysis of trade policy changes often suggest much higher elasticities (see Ruhl (2008) for a detailed discussion). To match sector-level trade balance dynamics closely, we set the goods Armington elasticity to 3 and the services elasticity to 1. We then use equilibrium conditions (marginal product pricing and zero profits) to calibrate μ_j^{us} from the input-output matrix. The scale factors M_j^{us} follow immediately. We use equilibrium conditions in a similar procedure to calibrate the investment sector's parameters.

Household and government parameters

We set the elasticity of intertemporal substitution $1/(1-\psi)$ to 0.5. We set the long-run interest rate to 3 percent. We set the discount factor β so that this interest rate is consistent with balanced growth. We set the elasticity of substitution between goods and services in consumption, $1/(1-\rho)$, to 0.5, as is Stockman and Tesar (1995). The household's first-order conditions imply

(25) $$\frac{\varepsilon^{ush}}{1-\varepsilon^{ush}} = \frac{p_{g1992}^{us}}{p_{s1992}^{usS}}\left(\frac{c_{g1992}^{ush}}{c_{c1992}^{ush}}\right)^{1-\rho}$$

which we use to calibrate the share parameter ε^{ush}. We then use data on hours worked to calibrate η. A similar procedure yields the government's share parameter ε^{usg}.

U.S. initial conditions

To calculate the initial capital stock we set the 1992 real interest rate to 4 percent. The real interest rate in 1992 is not an equilibrium object in our model; it would be determined in 1991 and 1992 is our initial year. The real interest rate on 10-year U.S. treasury bonds is approximately 4 percent in 1992. Our results are not sensitive to alternative approaches to calibrating the initial capital stock. Depreciation was 11.7 percent of GDP in 1992, so given a tax rate on capital income τ_k^{us} we can then calibrate the initial capital stock as

(26) $$k_{1992}^{us} = \frac{(1-\tau_k^{us})r_{k1992}^{us}k_{1992}^{us} - \delta k_{1992}^{us}}{r_{1992}}.$$

and the depreciation rate as

(27) $$\delta = \frac{11.7}{k_{1992}^{us}}.$$

We choose $\tau_k^{us} = 0.415$, which implies a value of $\delta = 0.062$, well within the standard range of annual depreciation rates used in the literature. U.S. government debt was 42.8 percent of GDP in 1992. We use this figure to set the government's bonds in 1992, $\bar{b}_{1992}^{usg}$, then set private bond holdings, $\bar{b}_{1992}^{ush}$, so that total net foreign assets $\bar{b}_{1992}^{ush} + \bar{b}_{1992}^{usg}$ are –6 percent of GDP as reported in Lane and Milesi-Feretti (2007).

Calibrating the rest of the world

To calibrate the remaining parameters we need to specify what the "rest of the world" is in the data. We calculate the United States' top 20 trading partners, ranked by average annual bilateral trade (exports plus imports) between 1992 and 2011, and weight them by their average share of U.S. total annual trade (again, exports plus imports) during this period. We use these countries weights to construct a composite trading partner, thinking of the rest of the world as being composed of 20 identical countries that all look like this composite. To calculate the rest of the world's gross output of goods and services, we take a weighted average of goods and services output of these 20 countries and multiply these figures by 20 to get total output of goods and services in the rest of the world. We use equilibrium conditions in a similar manner as before to calibrate the rest of the world's Armington aggregators and preference parameters.

Exogenous processes

We use historical data and future projections from the United Nations World Population Prospects 2010 Revision to construct time series for the demographic parameters for both the United States and the rest of the world (using the same 20 countries and weights as before). We use the "medium" scenario for the future projections. The United States and the rest of the world are projected to grow at different rates well past the 100-year terminal date in our model, so to ensure balanced growth in our computation we assume that populations in both countries begin to converge to constant levels after 2050. Our model's equilibrium dynamics between 1992 and 2030, the period on which we focus, are not sensitive to this assumption.

We calculate sector-level productivity growth rates using data on value added and labor compensation by sector from the BEA for the period 1987–2011. We use this data to perform growth accounting by sector, and find that the average growth rates of labor productivity over this period are 4.3 percent in goods, 1.3 percent in services, and -0.04 percent in construction. We use these values in the model between 1992 and 2030, and in the years following, we assume that all of the sector-level growth rates converge slowly to 2 percent per year, to ensure that the equilibrium converges to a balanced growth path.

We construct several time series for government consumption expenditure and debt. We use historical data from the NIPA tables for government consumption expenditures and from the U.S. Congressional Budget Office (CBO) for government debt. We use CBO projections as a starting point for our own projections, but make adjustments to allow for balanced growth in the long run. The CBO's long-run projections for government debt as a fraction of GDP vary greatly from year to year and are, quite frankly, implausible. The 2012 Long Term Budget Outlook provides two possible scenarios for government debt as a fraction of GDP in the long run, neither of which are stationary. The first, the "extended baseline" scenario, has debt as a fraction of GDP falling below zero in the long run. The other, the "extended alternative fiscal scenario," has government debt reaching more than 250 percent of GDP by 2045. These two scenarios differ substantially from the 2011 projections.

In the no-savings glut counterfactual scenario, we assume that government spending as a fraction of GDP remains constant, and that government debt gradually rises to 60 percent of GDP over time. In modeling the savings glut, we use the actual data for 1992–2011. In the gradual rebalancing scenario, consumption spending as a fraction of GDP rises gradually to approximately 23 percent in the long run, while the debt-GDP ratio converges to 74 percent. In the sudden stop scenario, we use the same spending series but assume that the debt-GDP ratio falls to 60 percent once the savings glut ends, reflecting a permanent change in U.S. government policy that coincides with the sudden stop.

6. Quantitative results

We begin this section by discussing our baseline model's historical performance in matching the U.S. trade balance (aggregate and disaggregated), real exchange rate, and sectoral labor dynamics between 1992 and 2012, and comparing these results to the counterfactual scenario in which the savings glut never took place. We then study the model's predictions for the future and how a sudden stop in foreign lending in 2015–2016 would play out. We finish with a discussion of the welfare consequences of the savings glut and a sudden stop.

Key facts and gradual rebalancing versus the no-savings glut counterfactual

As figure 6 shows, we calibrate our model to match the U.S. trade balance exactly between 1992 and 2012 exactly. We therefore cannot judge our model's performance in matching this fact, so we treat it as "fact zero." In this figure we have also plotted the trade balance from the model in which the savings glut never occurs. Notice that in this counterfactual scenario, trade would be close to balanced throughout the period. In the baseline model, however, the savings glut has occurred, so the U.S. must repay its debt in the long run. The figure shows that, as long as the rebalancing process goes smoothly, the U.S. trade balance will switch from a deficit to a surplus in 2018, and will reach a surplus of more than 0.8 percent of GDP by 2024.

Figure 7 plots the model's real exchange rate in the baseline model against the data and the no-savings glut counterfactual. Our model does a good job of matching the magnitude of the appreciation during the 1992–2012 period; the real exchange rate falls by 21.8 percent in the data and 21.7 percent in the model before beginning to depreciate. As we have mentioned, however, we fail to get the timing right. In the data, the real exchange rate begins to depreciate in 2002, 4 years before the trade deficit begins to shrink. In our model, the real exchange rate moves in tandem with the trade balance, so it does not begin to depreciate until 2006. In our model, consumers and firms begin to import fewer foreign goods and services once they begin to become more expensive. Perhaps the standard explanations in the J-curve literature (Backus, Kehoe, and Kydland, 1994) — time-to-build and import pattern adjustment frictions — can account for this fact, but we believe that these stories alone are not likely to explain the substantial 4-year lag. Another possible explanation for this pattern is the increase in importance of China in U.S. trade during the period. In figure 10 we decompose the U.S. real exchange rate into the real exchange rates with China and 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 much like the real exchange with China. To accurately capture this we would need to model a world with (at least) 3 countries, which is outside the scope of this paper. Figure 7 also shows that if the savings glut had never occurred at all, the U.S. real exchange rate would depreciate slowly over the long run, due in large part to differences in demographic changes as compared to the rest of the world. The savings glut did occur, however, and because the United States must run a trade surplus in the long run, its goods and services must become cheaper. Our model predicts that the real exchange rate will continue to appreciate for several years, eventually converging to a level that is 7 percent higher than that in the no-savings glut counterfactual.

Figure 8 plots the sector-level trade balances in the data and the model. We match both the goods and services trade balances closely between 1992 and 2012; in fact we match the goods trade balance almost exactly through the period. This aspect of the model's performance is due to our choice of Armington elasticities. Had we used the same elasticity in both sectors, the goods trade balance would not have moved enough, while the services trade balance would have been far too volatile. The figure also shows that in the absence of the savings glut, the services trade balance would not have been substantially different, while the goods trade balance, reflecting the aggregate trade balance, would have been almost flat. In the long run, the model predicts that the goods trade balance will be 1.4 percent of GDP higher by 2024 than it would have been if the savings glut had never occurred. Despite this, the model predicts that the goods trade balance will remain negative in perpetuity. In other words, the services sector will be the source of the entire long-run trade surplus. In the figure we see that the services sector's trade surplus will be only slightly higher in the long run than it would have been had the savings glut never occurred; The United States is prone to running a trade surplus of approximately 1 percent of GDP regardless of the circumstances.

Trade deficits and employment

Figure 9 plots the employment shares for goods and services in the model and the data. In the baseline gradual rebalancing scenario our model matches the data closely between 1992 and 2001, after which the goods employment share falls more in the data than in the model. For the 1992–2011 period, our baseline model captures 71 percent of the decline in the goods sector's employment share that we observe in the data. In the no-savings glut counterfactual, however, the model still captures 58 percent of the decline (the decline in the counterfactual is more than 80 percent of the decline in the baseline rebalancing scenario). The figure suggests that the savings glut did temporarily accelerate the decline in goods-sector employment, but on the whole our results indicate that the savings glut has played only a small role in the decline in the goods sector's employment share. The bulk of the decline is attributable to the simple fact that labor productivity has grown substantially faster in the goods sector than in the rest of the economy — the number of workers required to meet demand for U.S. goods has fallen much faster than the number of workers required to meet demand for U.S. services or construction.

The implication of this result is that, contrary to popular belief, the end of the savings glut will have very little impact on employment in the goods sector. By 2024, the model indicates that the goods sector's share of employment will be almost exactly the same regardless of whether or not the savings glut had happened. Despite the fact that the U.S. economy must run a trade surplus in the long run to repay its debt, the goods sector's employment share will continue to decline once the savings glut ends. There are two main reasons for this result. First, labor productivity continues to grow faster in goods than in other sectors in our model even after the savings glut ends. This aspect of structural change has been a consistent force in the U.S. economy since the 1960s, and we see no reason that it should end when the savings glut does. Second, the U.S. can repay its debt with the trade surplus generated by the services sector; our model predicts that the U.S. will continue to run a deficit in goods trade once the savings glut ends and the aggregate U.S. trade balance reverses.

Figure 9 also shows that the model captures several aspects of the construction sector's employment share between 1992 and 2011. Between 1992 and 2006, the construction sector's employment share rises in both the model and the data, although our model generates a larger increase. In other words, our model over-explains the boom in construction sector employment just by introducing the savings glut. The subsequent bust is smaller in the model than in the data, but this primarily due to the fact that 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 additional features to the model to more accurately model the crisis we would undoubtedly do better in this regard, but this is not the focus of our paper. As the figure shows, our model suggests that the effects of the savings glut on the construction sector will largely dissipate by 2016; just like in the goods sector, the construction sector's long-run employment share dynamics will be driven primarily by productivity growth rather than the rebalancing process.

One aspect of the data that our model fails to capture is the U.S. real interest rate. Figure 11 shows that in the data, the U.S. real interest rate (here measured as the ex-post real return on 10-year U.S. treasury bonds) fell by roughly 400 basis points between 1992 and 2012 (ignoring the temporary spike caused by deflation in 2009). By contrast, the real interest rate in our model moves very little, and the interest rate in the gradual rebalancing scenario is similar to that in the no-savings glut counterfactual; in our model the savings glut has little impact on the U.S. real interest rate. In our numerical experiments, we can obtain large real interest rate movements only by setting several elasticities of substitution (between goods and services in consumption and between domestic and foreign inputs) to extremely high values, which makes all other variables behave wildly. In other words, trade deficits can affect the U.S. real interest rate substantially in a model in which foreign consumption is perfectly substitutable for U.S.

consumption (a one-good model), but such a model has no hope of accurately capturing other important features of the data.

A sudden stop in 2015–2016

The results we have presented so far about the impact of the end of the savings glut on the U.S. economy focus on the scenario in which the rebalancing process is gradual and orderly. Figures 10–12 illustrate what our model predicts if this process instead involves a sudden stop in 2015–2016. In figure 12, we plot the trade balance and the real exchange rate in our model in both the gradual rebalancing and sudden stop scenarios. We see that a sudden stop will trigger an immediate reversal in capital flows and a large real exchange rate depreciation. The trade balance will rise from -1.6 percent of GDP to 2.7 percent on impact, and the real exchange rate will rise by 13 percent. We also see, however, that these effects are short-lived. By 2024 the trade balance and real exchange rate are almost identical regardless of whether a sudden stop happens or not. While a sudden stop involves quicker debt repayment, the effect on the long–term need to repay is small in our model. The United States has borrowed so much from the rest of the world in the last two decades that two years of rapid repayment simply don't make much of a dent in the overall stock of debt.

Figure 13 illustrates that the large, temporary trade balance reversal comes primary from the goods sector. The goods trade balance rises by 3.6 percent of GDP on impact, while the services trade balance rises by 0.6 percent. Figure 14 shows that despite the large increase in the goods trade balance, we do not see a large reallocation of labor towards the goods sector. The goods sector's labor compensation share rises from 14.8 percent to 15.5 percent on impact, but this is short-lived, just like the other effects of the sudden stop. Once the sudden stop ends, rapid productivity growth continues to push goods employment downward, and by 2024 the goods sector's labor compensation share is almost exactly the same regardless of whether the sudden stop occurs or not. As compared to the goods sector, construction sees a much larger (although just as temporary) impact from the sudden stop. The construction sector's share of labor compensation falls by more than one third, from 6.2 percent to 4.1 percent. There are two channels at work. First, construction is the only nontraded sector, so it is more sensitive to changes in capital flows (the same mechanism as in Kehoe and Ruhl, 2009). The second is that, as we have explained in discussing our calibration, the bulk of the construction sector's output is used to produce investment. The sudden stop causes the U.S. real interest rate to rise by more than 2.5 percent, leading to a large decline in investment, which substantially reduces demand for construction.

In short, a sudden stop in the United States will, in many respects, look similar to past sudden stop episodes in emerging economies. There will be a large trade balance reversal and real exchange rate depreciation, and severe disruption in the construction sector, but these effects will be short-lived — once the sudden stop ends, the U.S. economy will return to roughly the same trajectory on which it would have been had the sudden stop never happened. While a sudden stop will temporarily stem the decline in goods-sector employment, this trend will continue unabated once the sudden stop ends.

Welfare implications of the savings glut

We have shown that the savings glut has had a large impact on some aspects of the U.S. economy, and that a sudden stop will cause substantial short-term disruption. Here we ask two closely related questions: Did the savings glut make U.S. households better or worse off? How does the answer depend on whether the savings glut ends in a gradual rebalancing or a sudden stop?

To answer these questions, we construct a measure of real income in 1992 denominated in 1992 U.S. dollars. Our real income index is based on an alternative — but equivalent — specification of the representative household's preferences that is homogeneous of degree one,

$$(28) \qquad \left(\sum_{t=1992}^{\infty} \beta^t \left(\varepsilon^{us} \left(\frac{c_{gt}^{ush}}{n_t^{us}} \right)^{\rho} + (1-\varepsilon^{us}) \left(\frac{c_{st}^{ush}}{n_t^{us}} \right)^{\rho} \right)^{\frac{\eta\psi}{\rho}} \left(\frac{\bar{\ell}_t^{us} - \ell_t^{us}}{\bar{\ell}_t^{us}} \right)^{(1-\eta)\psi} \right)^{\frac{1}{\psi}} .$$

The cost of achieving this utility in units of the U.S. CPI in 1992 is

$$(29) \qquad \sum_{t=1992}^{\infty} \left(\prod_{s=1992}^{t-1} q_s \right) \left(p_{gt}^{us} c_{gt}^{ush} + p_{st}^{us} c_{st}^{ush} + w_t^{us} (\bar{\ell}_t^{us} - \ell_t^{us}) \right),$$

which we can write as

$$(30) \qquad \sum_{t=1992}^{\infty} \left(\prod_{s=1992}^{t-1} q_s \right) w_t^{us} \bar{\ell}_t^{us} + r_{k1992}^{us} \bar{k}_{1992}^{us} + p(p_{g1992}^{us}, p_{s1992}^{us})(\bar{b}_{1992}^{ush} + \bar{b}_{1992}^{usg}) .$$

Notice that U.S. households are ultimately responsible for their government's debt. The prices and quantities above represent equilibrium objects in our benchmark gradual rebalancing scenario, the one in which the savings glut occurs, but a sudden stop does not. To convert this object to 1992 dollars, we scale it so that consumption expenditures in the model are equal to 1992 private consumption in the NIPA tables, C_{1992}^{ush}. The scaling factor,

$$C = \frac{C_{1992}^{ush}}{p_{g1992}^{us} c_{g1992}^{ush} + p_{s1992}^{us} c_{s1992}^{ush}}, \tag{31}$$

converts consumption expenditures in the model into 1992 dollars, and

$$P = \frac{\sum_{t=1992}^{\infty} \left(\prod_{s=1992}^{t-1} q_s \right) w_t^{us} \bar{\ell}_t^{us} + r_{k1992}^{us} \bar{k}_{1992}^{us} + p(p_{g1992}^{us}, p_{s1992}^{us})(\bar{b}_{1992}^{ush} + \bar{b}_{1992}^{usg})}{\left(\sum_{t=1992}^{\infty} \beta^t \left(\varepsilon^{us} \left(\frac{c_{gt}^{ush}}{n_t^{us}} \right)^{\rho} + (1-\varepsilon^{us}) \left(\frac{c_{st}^{ush}}{n_t^{us}} \right)^{\rho} \right)^{\frac{\eta\psi}{\rho}} \left(\frac{\bar{\ell}_t^{us} - \ell_t^{us}}{\bar{\ell}_t^{us}} \right)^{(1-\eta)\psi} \right)^{\frac{1}{\psi}}} \tag{32}$$

converts utility into expenditures, where P is the price of utility. The real income of U.S. households in1992 dollars in our benchmark scenario is

$$PC \left(\sum_{t=1992}^{\infty} \beta^t \left(\varepsilon^{us} \left(\frac{c_{gt}^{ush}}{n_t^{us}} \right)^{\rho} + (1-\varepsilon^{us}) \left(\frac{c_{st}^{ush}}{n_t^{us}} \right)^{\rho} \right)^{\frac{\eta\psi}{\rho}} \left(\frac{\bar{\ell}_t^{us} - \ell_t^{us}}{\bar{\ell}_t^{us}} \right)^{(1-\eta)\psi} \right)^{\frac{1}{\psi}}. \tag{33}$$

To calculate real income in alternative scenarios, like the counterfactual in which the savings glut does not occur or the scenario in which both the savings glut and sudden stop occur, we simply multiply lifetime utility by the scaling factors obtained from the benchmark scenario:

$$PC \left(\sum_{t=1992}^{\infty} \beta^t \left(\varepsilon^{us} \left(\frac{\tilde{c}_{gt}^{ush}}{n_t^{us}} \right)^{\rho} + (1-\varepsilon^{us}) \left(\frac{\tilde{c}_{st}^{ush}}{n_t^{us}} \right)^{\rho} \right)^{\frac{\eta\psi}{\rho}} \left(\frac{\bar{\ell}_t^{us} - \tilde{\ell}_t^{us}}{\bar{\ell}_t^{us}} \right)^{(1-\eta)\psi} \right)^{\frac{1}{\psi}} \tag{34}$$

where use tildes to denote equilibrium objects in the alternative scenario being studied.

In our baseline model we assume that in 1992, model agents expect government consumption expenditures to remain fixed at the 1992 level of 16.6 percent of GDP, but when the savings glut begins an unforeseen change in government spending policy occurs: government spending as a fraction of GDP tracks the data between 1993 and 2011, then rises to 22.9 percent over time. This reflects policy changes that have occurred over the past two decades, e.g. increased healthcare and defense spending, that people likely did not anticipate in the early

1990s. This increase in government consumption gives U.S. households an incentive to save for the future. We report our welfare results under the alternative assumption: in 1992, agents expect government consumption as a fraction of GDP to follow the path it actually took between 1992 and 2011, and then follow the same trajectory to 22.9 percent over time that we used in the savings glut scenario in our main exercise. In the savings glut and sudden stop scenario, we require that government consumption, in terms of actual quantities of goods and services, stay constant in all three stages of the exercise, the no savings glut counterfactual, the savings glut with gradual rebalancing, and the sudden stop. This modification has virtually no impact on any of the results reported above, and it allows for direct welfare comparisons across the three scenarios even if government spending enters the utility function — as long as it enters in an additively separable fashion — allowing us to ask whether the savings glut is good or bad, and just how costly a sudden stop would be.

The first column of table 3 presents our results on the welfare impact of the savings glut and sudden stop for our baseline model. In panel (a) we report real income of U.S. households in 1992 relative to the gradual rebalancing scenario. We see that real income of U.S. households in 1992 would have been $679 billion lower if the savings glut had never occurred. This is a large sum, more than 10 percent of US GDP in 1992. Twenty years of increased consumption of foreign imports have made U.S. households substantially better off — as long as the savings glut ends in gradual rebalancing. If the savings glut ends in a sudden stop, however, 1992 real income of U.S. households will fall by $1,034 billion compared to the gradual rebalancing scenario. In other words, if the savings glut ends in a sudden stop rather than gradual rebalancing, U.S. households would have been better off if the savings glut had never occurred at all, despite the fact that the welfare impact of the sudden stop is discounted by 23 years (1992 to 2015). To see how much of this result is driven by our assumption that a sudden stop is accompanied by a TFP shock, we also report welfare in a version of our model in which a sudden stop occurs but TFP does not fall. In this case, welfare falls by only $390 billion — a sudden stop without an accompanying TFP shock will be painful, but will not completely wipe out the welfare gains generated by the savings glut. In panel (b), we perform a similar exercise from the perspective of U.S. households in 2015 to see how painful is the sudden stop in the year it occurs, eliminating the impact of discounting. With a TFP shock, a sudden stop will reduce real income in 2015 by $3,588 billion. Even without the TFP shock, real income falls by $1,353 trillion.

7. Sensitivity analysis

We have performed a wide range of sensitivity analyses with our model, like adding factor adjustment costs, varying elasticities of substitution, changing our assumptions about government spending and debt, and reducing the size of the rest of the world's economy relative to that of the United States. Our main results about reallocation of labor across sectors and the impact of a sudden stop are quite robust to all of most of these modifications, so in the interest of brevity we have chosen not to discuss them here. There are four extensions, however, that we wish to discuss in detail. First, we relax our baseline model's assumption of perfect foresight by studying a version of our model in which the duration of the savings glut is uncertain — ω_t^{rw} now follows a stochastic process — and agents have rational expectations. Second, we study how our results change when re remove the input-output production structure. Third, we study a version of our model in which services cannot be traded internationally. Last, we study a version of our model in which U.S. borrowing is driven by the preferences of U.S. households rather than households in the rest of the world — a "domestic savings drought" instead of a global savings glut.

Stochastic savings glut

In our baseline model, agents in both the United States and the rest of the world have perfect foresight once the savings glut begins; they know exactly when it will end and the rate at which it will rebalance. We have also studied a version of our model in which there is uncertainty about the length of the savings glut. In this version of the model, once the savings glut begins there is a 10 percent chance in each year between 1993 and 2011 that the savings glut will end in the following period, and the rest of the world's demand for saving will begin to increase again. The other 90 percent of the time, the savings glut will continue for at least one additional period. The realized path the economy takes is the one in which the savings glut persists through 2011, and while this is the unconditionally most likely path the economy can take, it is not very likely from the perspective of model agents in 1992. Our results with this version of the model indicate that this kind of uncertainty has no discernible impact on our results. We do not plot the results from the model with uncertainty, because they are virtually indistinguishable from our baseline results.

This addition to our model represents a substantial technical contribution. Due to the presence of asymmetric, time-varying growth rates in productivity, demographics, and other variables, our modeling framework does not admit a stationary dynamic program. In the model with uncertainty, the current value of the stochastic savings glut process is not a sufficient statistic for the exogenous state of the economy — the entire history of shocks matters. As a consequence, we must solve for the growth paths of the world economy along all possible sequences of shocks simultaneously. The number of possible sequences increases in proportion to the number of periods with uncertainty, so the dimensionality of the problem increases rapidly. To our knowledge, no other studies have attempted to solve this kind of model. We believe this framework was a wide variety of applications.

No input-output production structure

Our baseline model incorporates an input-output structure in which each sector of the U.S. economy purchases intermediate inputs from other sectors, and combines these inputs with capital and labor to produce output. This is not a typical ingredient in most models in the international macroeconomics literature, although some studies such as Kehoe and Ruhl (2009) take a similar approach. Here we demonstrate the importance of this approach for our model's results by studying an alternative version of our model in which we remove this input-output structure. We set all intermediate input values in our 1992 input-output matrix to zero, then adjust the remainder of the matrix so that it is once again consistent with the national accounts and sectoral labor compensation data for 1992. We then re-calibrate our model and perform the same quantitative exercises described above: gradual rebalancing, the no-savings glut counterfactual, and the sudden stop.

Removing the input-output structure has little impact on our model's predictions for the trade balance (aggregate and disaggregate) and the real exchange rate, so we focus on the employment results. Figure 15 plots the goods sector's labor compensation share in the model without intermediate inputs against the baseline model's results and the data. The first thing we wish to point out is that in the no-savings glut counterfactual, the goods sector's labor compensation share declines at a much lower rate than in the baseline model. By 2024, that share falls to 16.6%, versus 13.0% in the baseline model. In other words, removing intermediate inputs from the model reduces the amount of structural change in the model substantially. The structural change literature emphasizes the importance of the elasticity of substitution between goods and services in consumption for reallocation of labor across sectors that is driven by asymmetric productivity growth; a lower elasticity leads to more reallocation (see, for example, Ngai and Pissarides, 2007). In our calibration, that elasticity is 0.5, a standard value in the literature. The Leontief input-output structure we employ implies that the elasticity of substitution between goods and services in intermediate usage is, and as a consequence the elasticity of substitution between goods and services in gross output is less than 0.5. When we remove intermediates, the elasticity in gross output rises substantially, lowering the amount of reallocation of labor from goods to services in the long run.

The second point we wish to make is that removing intermediate inputs leads our model to attribute a much larger decline in goods-sector employment to the savings glut than does our baseline model — albeit a temporary one, just as before. The path of the goods sector's share of labor compensation between 1992 and 2006 in the model without intermediates is very similar in the path ion the baseline model. Relative to the no-savings glut counterfactual, however, the 1992–2006 decline in goods-sector employment is larger in the model without intermediates than the baseline model. In the no-intermediates model, the goods sector's labor compensation share in 2006 is 15.3% in the gradual rebalancing scenario compared to 17.3% in the no-savings glut counterfactual, a difference of 2%. In the baseline model, that share is 14.9% in the rebalancing scenario versus 16.5%, a difference of only 1.4%. The reverse is also true — the end of the savings glut leads to a larger long-run increase in goods-sector employment relative to the no-savings glut counterfactual. In the model without intermediates, the goods sector's labor compensation share in 2024 is 16.6% in the no-savings glut counterfactual versus 17.5% rebalancing scenario. In the baseline model, those numbers are closer, 13.0% and 13.6%. In short, removing intermediates leads our model to overstate the effects of both the savings glut and its end on goods-sector employment.

Last, our model predicts that a sudden stop will have a larger (albeit temporary) impact on goods-sector employment when we remove intermediate inputs. In the no-intermediates model, the goods sector's labor compensation share rises from 16.9% in 2014 to 19.2% in 2015 with a sudden stop, compared to an increase from 14.9% to 15.2% in the baseline model. The reason for this is that goods are used more as intermediates than services, and the TFP shock that accompanies the sudden stop reduces intermediate demand. This channel disappears when we remove intermediates from the model. So removing intermediates from the model not only causes us to overstate the impact of the savings glut on goods-sector employment, it also causes us to overstate the short and long-run impact of the end of the savings glut on goods-sector employment in both the gradual rebalancing and sudden stop scenarios.

Nontradable services

Modelers in the international macroeconomics literature typically assume that services output is nontradable. By contrast, in our baseline model, both goods and services can be traded internationally, and we have calibrated our model to reflect two observations about the data: (i) the United States typically runs a trade surplus in services and a deficit in goods; and (ii) the U.S. goods trade balance is more volatile than the U.S. services trade balance. Here, we study how our results change when we adopt the standard modeling convention in which services and nontradable. We re-calibrate our model as in the previous section so that the goods sector is responsible for total imports and exports in the 1992 national accounts, and then perform the same quantitative exercises described above.

This change has almost no impact on any of our results except for goods-sector employment and welfare. Figure 17 plots the goods sector's labor compensation share in the model with nontradable services against the baseline model's results and the data. Making services nontradable has almost no impact on the model's predictions for the no-savings glut counterfactual. In the gradual rebalancing scenario, making services nontradable increases the impact of the savings glut on goods-sector employment during the savings glut itself, but this effect is only temporary. In the model with nontradable services, the goods sector's labor compensation share falls from 19.7% in 1992 to 14.5% in 2006, versus 14.9% in the baseline model, an additional cumulative drop of 0.4%. This cumulative difference begins to shrink after 2006, however, and by 2024 the goods sector's labor compensation share in the model with nontradable services is almost exactly the same as in the baseline model. The figure also shows that in the model with nontradable services, a sudden stop in 2015–2016 will cause a larger reallocation of labor into the goods sector that in the baseline model, but this effect is also temporary. Finally, in the third column of table 3 we see that making services nontradable reduces the increase in 1992 real income generated by the savings glut from $679 billion to $482 billion, a difference of 30%.

By reducing the price of imported goods and services through real exchange rate appreciation, the savings glut causes U.S. consumers and producers to substitute foreign goods and services for domestic ones. This leads to a reduction in the share of labor compensation paid in tradable sectors. When we assume services are nontradable, the goods sector must bear the full brunt of this reduction; ignoring the tradeability of U.S. services leads would lead us to overstate the impact of the savings glut on goods-sector employment. As we have shown, however, the labor-market effects of the savings glut are largely temporary — by 2024 the goods sector's share of total labor compensation is roughly the same as it would have been had the savings glut never happened at all. As a consequence, ignoring the tradeability of U.S. services has little impact on our model's predictions for the long-run path of goods-sector employment.

Domestic savings drought

In our baseline model, we have adopted the global savings glut hypothesis proposed by Bernanke (2005), which states that U.S. borrowing from the rest of the world since the early 1990s has been driven primarily by increased demand for saving in the rest of the world. A number of studies, such as Obstfeld and Rogoff (2009), Chinn and Ito (2007), and Gruber and Kamin (2007), argue that domestic factors such as monetary policy, housing market policy, and innovations in financial markets were the primary cause of U.S. borrowing. In this section, we study a version of our model in which the preferences of U.S. households, rather than households in the rest of the world, that drives the U.S. trade balance. We follow Chinn and Ito (2007) and refer to this version of the model as the "domestic savings drought" model. In the savings drought model, the preferences of U.S. households take a the same form as in (12). As before, we calibrate the preference parameter ω_t^{us} so that the model matches the U.S. trade balance exactly during 1992–2012, after which it gradually converges to its long-run level of 1. Because investment dynamics are our focus in this exercise, we add frictions to adjusting capital (in both the savings glut and savings drought models) following Lucas and Prescott (1971) to allow our model to more closely match investment volatility we see in the data. This extension has very little impact on our main results.

To assess which of the models is more consistent with the data, we focus on investment and construction employment. If U.S. borrowing has been driven by increased demand for saving in the rest of the world (a global savings glut), U.S. households should have taken advantage of cheap foreign goods to increase both investment and consumption since the relative value they placed on future consumption remained unchanged. Since construction is an important input to the production of investment, this should have led to an increase in the construction sector's employment share. As we will see, both the model and the data are consistent with this story. If, on the other hand, U.S. borrowing has been driven by reduced demand for saving at home, U.S. investment and construction employment should have fallen as a result of decreased preference for future consumption. This is not what we see in the data. In short, the global savings glut hypothesis suggests that investment and construction employment should move in the opposite direction as the trade balance, while the domestic savings drought hypothesis suggests that they should move in the same direction. The former is true in the data.

Figure 17 plots investment as a share of GDP in the baseline model and the savings drought model against the data. Except for in 1993, investment in the baseline model moves in the same direction as the data. By contrast, investment in the savings drought model falls dramatically beginning in 1997 while it continues to rise in the data (except for during the 2001 recession which we have not attempted to incorporate at all). During the financial crisis of 2008–2009 (which we have modeled solely through the increasing trade balance), investment falls in the baseline model and the data, but rises in the savings drought model. The correlation between investment in the model and the data in first differences is 0.72 for the baseline model and -0.63 in the savings drought. Figure 18 plots the construction sector's labor compensation share in the two models and the data, and the same pattern emerges. In the baseline model, the construction employment share generally moves in the same direction as the data, whereas in the savings drought model, the construction employment share begins to fall in 1997. The correlations in first differences are 0.57 for the baseline model and -0.44 for the savings drought model.

Clearly, the baseline model is more consistent with the data on investment and construction employment. As suggested above, we can also look at the correlations (again in first differences) between these two variables and the trade balance. In the data, the correlation between investment and the trade balance is -0.79. In the baseline model, this correlation is -0.80, whereas it is 0.83 in the savings drought model. The correlation between the construction sector's labor compensation share and the trade balance is -0.60 in the data, -0.82 in the baseline model, and 0.85 in the savings drought model. These results indicate that our models capture well the intuitive arguments described above about the implications of a global savings glut versus a domestic savings drought for investment and construction employment, and that the data is consistent with the implications of a savings glut, not a savings drought.

To sum up, while our baseline model's results are consistent with the data on investment and construction employment, the "savings drought" version of our model is not; changes in investment and the construction sector's labor compensation share in the savings drought model are significantly and negatively correlated with the data. This finding serves as both a justification of our approach to modeling U.S. borrowing and evidence in favor of the savings glut hypothesis.

8. Conclusion

This paper studies the impact of the global savings glut — an increased willingness on the part of the rest of the world to trade their own goods and services in the present for claims on U.S. goods and services in the future — on the U.S. economy over the past two decades as well as the next one. We build a model of the U.S. and the rest of the world that incorporates a number of unique features and show that it accounts for four key facts about the U.S. economy during 1992–2011. First, the trade deficit increased then decreased. Second, the real exchange rate appreciated then depreciated roughly at the same time. Third, the trade balance dynamics are driven almost entirely by the goods trade balance. Finally, labor shifted away from the goods sector towards services and construction. We use our model to show that while faster productivity growth in the goods sector, compared to other sectors, is responsible for the bulk of the shift in employment away from the goods sector, the savings glut is in fact responsible for a large fraction of the boom in construction employment during this period.

We then use our model to ask what will happen in the future when the forces driving the savings glut begin to taper off, causing the United States to begin to pay back the debt it has incurred. We show that the United States will run a perpetual trade surplus, but will nevertheless continue to import more goods than it exports — this trade surplus will originate entirely in the services sector. The real exchange rate will continue to depreciate in the long run as the supply of foreign goods and services in the United States falls, and traditional structural change forces will continue to drive a decline in goods sector employment despite the fact that the United States will produce more of the goods it consumes at home. We also use our model to ask what will happen if the savings glut ends swiftly and unexpectedly, in a sudden stop episode like those that hit Mexico and Southeast Asian in the 1990s. We show that this scenario will trigger sharp increases in the trade balance and the real exchange rate and potentially painful reallocations of labor across sectors, but it will have little lasting impact — the long run trajectory of the U.S. economy will be approximately the same by 2024 regardless of whether a sudden stop occurs.

Our study identifies several puzzles. In our model, the trade balance and real exchange rate move simultaneously; an increase in the trade deficit is always accompanied by a real exchange rate appreciation (and the reverse for a decrease in the deficit). In the data, real exchange rate appreciation peaks in 2002, 4 years before the trade deficit peaks. We believe that modeling several of the United States' trade partners separately (China in particular) may help resolve this issue. Another puzzle our study identifies concerns the relationship between foreign lending and U.S. real interest rates. Contrary to popular wisdom, our results suggest that the savings glut is not an important factor in driving the low U.S. real interest rates of the past decade. Our results are not far off from empirical estimates on the effect of foreign lending on U.S. real interest rates (see, for example, Warnock and Warnock, 2008), indicating that researchers may need to look elsewhere. Nevertheless, we cannot rule out the possibility that the savings glut contributed to low U.S. real interest rates by interacting with domestic factors like those discussed by Bernanke, Bertaut, DeMarco, and Kamin (2011).

One of our main results is that the manner in which the savings glut ends — gradual rebalancing or sudden stop — will not have much impact on the long-run trajectory of the U.S. economy. We wish to leave the reader with one final point: the fact that the savings glut happened *does* have important implications for the future of the U.S. economy. The U.S. economy's current long-run trajectory is very different than the one it would have taken had the savings glut not taken place at all. Figure 6 illustrates this point by plotting the aggregate trade balance and real exchange rate in our gradual rebalancing scenario against the counterfactual in which the savings glut never happened. In the counterfactual, U.S. trade is approximately balanced in the long run since the United States has little debt to repay. Because the savings glut did happen, however, our model predicts that the U.S. will run a trade surplus of around one percent of GDP in perpetuity. So while the manner in which the savings glut ends is not likely to affect the U.S. economy in the long run, the fact that the savings glut occurred will have long-lasting effects.

References

Acemoglu, D. and V. Guerrieri (2006), "Capital Deepening and Non-Balanced Economic Growth," NBER Working Paper 12475.

Arellano, C. (2008), "Default Risk and Income Fluctuations in Emerging Economies," *American Economic Review*, 98, 690–712.

Backus, D., E. Henriksen, F. Lambert, and C. Telmer (2006), "Current Account Fact and Fiction," Stern School of Business, New York University, University of Oslo, and Tepper School of Business, Carnegie Mellon University.

Backus, D., P. J. Kehoe, and F. E. Kydland (1994), "Dynamics of the Trade Balance and the Terms of Trade: The J-Curve?," *American Economic Review*, 84, 84–103.

Baumol, W. J. (1967), "Macroeconomics of Unbalanced Growth: The Anatomy of Urban Crisis." *American Economic Review*, 57, 415–426.

Bernanke, B. S. (2005), "The Global Saving Glut and the U.S.Current Account Deficit," a 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.

Borio, C. E. V. and P. Disyatat (2011), "Global Imbalances and the Financial Crisis: Link or No Link?" BIS Working Paper 346.

Buera, F. J. and J. P. Kaboski (2009), "Can Traditional Theories of Structural Change Fit the Data?" *Journal of the European Economic Association*, 7, 467–477.

Buera, F. J. and J. P. Kaboski (2012), "Scale and the Origins of Structural Change," *Journal of Economic Theory*, 147(2), 684–712.

Buera, F. J. and J. P. Kaboski (2012), "The Rise of the Service Economy," *American Economic Review*, 102, 2540–2569.

Burstein, A., M. Eichenbaum, and S. Rebelo (2005), "Large Devaluations and the Real Exchange Rate," *Journal of Political Economy*, 113, 742–784.

Caballero, R.J., E. Farhi, and P.-O. Gourinchas (2008), "An Equilibrium Model of 'Global Imbalances' and Low Interest Rates," *American Economic Review*, 98, 358–393.

Caballero, R.J. and A. Krishnamurthy (2009), "Global Imbalances and Financial Fragility," *American Economic Review*, 99, 584–588

Calvo, G.A. (1988), "Servicing the Public Debt: The Role of Expectations," *American Economic Review*, 78, 647–661.

Calvo, G. A. (1998), "Capital Flows and Capital-Market Crises: The Simple Economics of Sudden Stops," *Journal of Applied Economics*, 1, 35–54.

Calvo, G. A., A. Izquierda, and L. F. Mejía (2004), "On the Empirics of Sudden Stops: The Relevance of Balance-Sheet Effects," NBER Working Paper 10520.

Calvo, G. A., A. Izquierdo, and E. Talvi (2006), "Phoenix Miracles in Emerging Markets: Recovering without Credit from Systemic Financial Crises," IADB Working Paper 570.

Calvo, G. A. and E. Talvi (2005), "Sudden Stops, Financial Factors and Economic Collapse in Latin America: Learning from Agentina and Chile," NBER Working Paper 11153.

Cecchetti, S.G., M.S. Mohanty, and F. Zampolli (2011), "The Real Effects of Debt," BIS Working Paper 352.

Chari, V. V., P. J. Kehoe, and E. R. McGrattan (2005), "Sudden Stops and Output Drops," *American Economic Review*, 95, 381–387.

Chinn, M. D. and H. Ito (2007), "Current Account Balances, Financial Development and Institutions: Assaying the World 'Savings Glut'," *Journal of International Money and Finance, 26, 546–569*.

Cole, H. L. and T. J. Kehoe (2000), "Self-Fulfilling Debt Crises," *Review of Economic Studies*, 67, 91–116.

Conesa, J. C., T. J. Kehoe, and K. J. Ruhl (2007), "Modeling Great Depressions: The Depression in Finland in the 1990s," Federal Reserve Bank of Minneapolis Quarterly Review, 31(1), 16–44.

Cook, D. E. and M. B. Devereux (2006), "External Currency Pricing and the East Asian Crisis," *Journal of International Economics*, 69, 37–63.

Dekle, R., J. Eaton, and S. Kortum (2007), "Unbalanced Trade," NBER Working Paper 13035.

Du, Q. and S. Wei (2012), "A Theory of the Competitive Saving Motive," Manuscript, Monash University and Columbia University

Duarte, M. and D. Restuccia (2010), "The Role of the Structural Transformation on Aggregate Productivity," *Quarterly Journal of Economics*, 125, 129–173.

Gourinchas, P.-O., and Rey, H. (2007), "From World Banker to World Venture Capitalist: US External Adjustment and the Exorbitant Privilege," in R.H. Clarida, ed., G7 *Current Account Imbalances: Sustainability and Adjustment*, University of Chicago Press, 11–66.

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

Greenspan, A. (2005), "Monetary Policy Report to the Congress," statement before the Committee on Banking, Housing, and Urban Affairs, U.S.Senate, February 16.

Hutchison, M. M., I. Noy, and L. Wang (2010), "Fiscal and Monetary Policies and the Cost of Sudden Stops," *Journal of International Money and Finance*, 29, 973–987.

Favilukis, J., S. C. Ludvigson, and S. Van Niewerburgh (2012), "Foreign Ownership of U.S. Safe Assets: Good or Bad?" London School of Economics and New York University

Frankel, J. A. and E. A. Cavallo (2004), "Does Openness to Trade Make Countries More Vulnerable to Sudden Stops, or Less? Using Gravity to Establish Causality," NBER Working Paper 10957.

Kaminsky, G. L. and C. M. Reinhart (1999), "The Twin Crises: The Causes of Banking and Balance-of-Payments Problems," *American Economic Review*, 89, 473–500.

Kehoe, P. J. and T. J. Kehoe (1994), "A Primer on Static Applied General Equilibrium Models," Federal Reserve Bank of Minneapolis Quarterly Review, 18, 2–16.

Kehoe, T. J. and K. J. Ruhl (2009), "Sudden Stops, Sectoral Reallocations and the Real Exchange Rate," *Journal of Development Economics*, 89, 235–249.

Kravis, I. B., A. W. Heston, and R. Summers (1983), "The Share of Services in Economic Growth," in F. G. Adams and B. G. Hickman, eds., *Global Econometrics: Essays in Honor of Lawrence R. Klein*, MIT Press, 188–218.

Krishnamurthy, A. and A. Vissing-Jorgensen (2007), "The Demand for Treasury Debt," NBER Working Paper 12881.

Krugman, P.R. (2007), "Will there be a Dollar Crisis?" *Economic Policy*, 22, 435–467.

Lucas, R.E. and E.C. Prescott (1971), "Investment Under Uncertainty," *Econometrica*, 39, 659–681.

McGrattan, E. R. and E. C. Prescott (2008), "Technology Capital and the U.S.Current Account," NBER Working Paper 13983.

Mendoza, E. G. (2001), "Credit, Prices and Crashes: Business Cycles with a Sudden Stop," in S. Edwards and J. Frankel, eds., *Preventing Currency Crises in Emerging Markets*, University of Chicago Press, 2002, 335–389.

Mendoza, E. G. (2005), "Real Exchange Rate Volatility and the Price of Nontradables in Sudden-Stop-Prone Economies," *Economia*, 6, 103–148.

Mendoza, E. G. (2010). "Sudden Stops, Financial Crises, and Leverage." *American Economic Review*, 100, 1941–66.

Mendoza, E. G., V. Quadrini, and J.-V. Rios-Rull (2009), "Financial Integration, Financial Development, and Global Imbalances," *Journal of Political Economy*, 117, 371–416

Meza, F. and E. Quintin (2007), "Factor Utilization and the Real Impact of Financial Crises," *B. E. Journal of Macroeconomics (Advances)*, 7, article 33.

Ngai, L. R. and Pissarides, C. A. (2007), "Structural Change in a Multisector Model of Growth." *American Economic Review*, 97, 429–443.

Obstfeld, M. (2012), "Does the Current Account Still Matter?" NBER Working Paper No. 17877.

Obstfeld, M. and Rogoff, K.S. (2007), "The Unsustainable U.S.Current Account Position Revisited," in R.H. Clarida, ed., *G7 Current Account Imbalances: Sustainability and Adjustment*, University of Chicago Press, 339–376.

Obstfeld, M. and K. Rogoff (2009), "Global Imbalances and the Financial Crisis: Products of Common Causes," CEPR Discussion Paper 7606.

Roubini, N. and Setser, B. (2005), "Will the Bretton Woods 2 Regime Unravel Soon? The Risk of a Hard Landing in 2005–2006," New York University and Oxford University

Sargent, T. J. (1978), "Estimation of Dynamic Labor Demand Schedules under Rational Expectations," *Journal of Political Economy*, 86, 1009–44.

Shin, H.S. (2012), "Global Banking Glut and Loan Risk Premium," *IMF Economic Review*, 60, 155–192.

Sposi, M. (2012), "Evolving Comparative Advantage, Structural Change, and the Composition of Trade," University of Iowa.

Song, Z. M., K. Storesletten, and F. Zilibotti (2010), "Growing like China," CEPR Discussion Papers 7149.

Stockman, A. C., and L. L Tesar (1995), "Tastes and Technology in a Two-Country Model of the Business Cycle: Explaining International Comovements," *American Economic Review*, 168–185

Summers, L. H. (2004), "The United States and the Global Adjustment Process," a speech at the third annual Stavros S. Niarchos Lecture, Institute for International Economics, Washington DC, March 23.

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 (Billions of 1992 Dollars)

	Inputs			Final demand					
Industry	**Goods**	**Services**	**Construction**	**Private consumption**	**Government consumption**	**Investment**	**Exports**	**-Imports**	**Total demand**
Goods	1,345	424	240	891	196	345	448	-545	**3,346**
Services	638	1,488	179	3,346	854	228	187	-123	**6,798**
Construction	26	139	1	-	-	514	-	-	**679**
Labor compensation	849	3,273	188	-	-	-	-	-	**4,310**
Returns to capital	488	1,474	71	-	-	-	-	-	**2,033**
Total gross output	**3,346**	**6,798**	**679**	**4,237**	**1,050**	**1,088**	**635**	**-668**	

Table 2: Calibration

Parameter	Value	Statistic	Target
Producer parameters			
A_g, A_s, A_c	(2.59,1.56,2.95)	Domestic gross output in 1992	(52.8,107,10.7)
a_{gg}, a_{sg}, a_{cg}	(0.41,0.19,0.01)	Share of intermediates in domestic goods in 1992	(0.41,0.19,0.01)
a_{gs}, a_{ss}, a_{cs}	(0.07,0.22,0.02)	Share of intermediates in domestic services in 1992	(0.07,0.22,0.02)
a_{gc}, a_{sc}, a_{cc}	(0.35,0.26,0.001)	Share of intermediates in construction in 1992	(0.35,0.26,0.001)
$\alpha_g, \alpha_s, \alpha_c$	(0.37,0.31,0.27)	Capital's share of domestic value added in goods/svcs/constr. in 1992	(0.37,0.31,0.27)
$\theta_g, \theta_s, \theta_c$	(0.32,0.21,0.47)	Share of intermediates in investment good production in 1992	(0.32,0.21,0.47)
G	2.85	Investment in 1992	17.2
Household parameters and initial conditions			
$\bar{b}^{ush}_{1992}$	36.8	Capital account balance in 1992, in percent of GDP	0.08
$\bar{k}^{us}_{1992}$	176.25	Real interest rate in 1992, in percent	4.00
β^{us}, β^{rw}	(0.996,0.996)	Long-term real interest rate, in percent	3.00
$\varepsilon^{ush}, \varepsilon^{rw}$	(0.07,0.19)	Goods share of private consumption in 1992	21.0
ρ	-1.00	Elasticity of substitution, traded to nontraded	0.50
η	0.29	Ratio of hours worked to available hours in 1992	0.29
ψ	-1.00	Intertemporal elasticity of substitution	0.50
δ	0.066	Depreciation to GDP in 1992, in percent	11.7
Trade parameters			
M^{us}_g, M^{us}_s	(1.78,1.08)	Gross composite output in goods/services in 1992	(61.3,109.1)
μ^{us}_g, μ^{us}_s	(0.65,0.98)	U.S.imports in 1992	(8.59,1.94)
M^{rw}_g, M^{rw}_s	(0.71,0.95)	Implied R.W. traded goods/services output	(86.9,161.7)
μ^{rw}_g, μ^{rw}_s	(0.71,0.98)	U.S.exports in 1992 in 1992	(7.06,2.95)
ζ_g, ζ_s	(0.67,0.00)	Elasticity of substitution, domestic traded to imports	(3.00, 1.00)
Government parameters and initial conditions			
$\bar{b}^{usg}_{1992}$	-42.8	U.S.government debt in 1992	42.8
τ^{us}_k	0.415	Depreciation rate	0.066
τ^{us}_{k1993}	0.397	Investment in 1992	17.2
ε^{usg}	0.179	Goods share of government consumption, in percent	0.19
Time series parameters			
$\{\tilde{n}^{us}_t, \tilde{n}^{rw}_t, \bar{\ell}^{us}_t, \bar{\ell}^{rw}_t\}^{\infty}_{t=0}$		U.N. Population Prospects: 2010 Revision	
$\{\omega^{rw}_t\}^{\infty}_{t=0}$		U.S.trade balance, 1992–2011	
$\{\gamma^{us}_{jt}\gamma^{rw}_{jt}\}^{\infty}_{t=0}$		Labor productivity growth in goods/svs/constr. 1987–2011	
$\{\nu_t, \upsilon_t\}^{\infty}_{t=0}$		CBO historical data and projections; authors' projections	

Table 3: Welfare impact of savings glut and sudden stop

Change in real income (billions of dollars)	No adjustment costs	Adjustment costs
(a) In 1992 compared to rebalancing scenario		
No savings glut counterfactual	-679	-821
Sudden stop (no TFP shock)	-390	-444
Sudden stop (TFP shock)	-1,034	-1,118
(b) In 2015 compared to rebalancing scenario		
Sudden stop (no TFP shock)	-1,353	-1,547
Sudden stop (TFP shock)	-3,588	-3,900

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

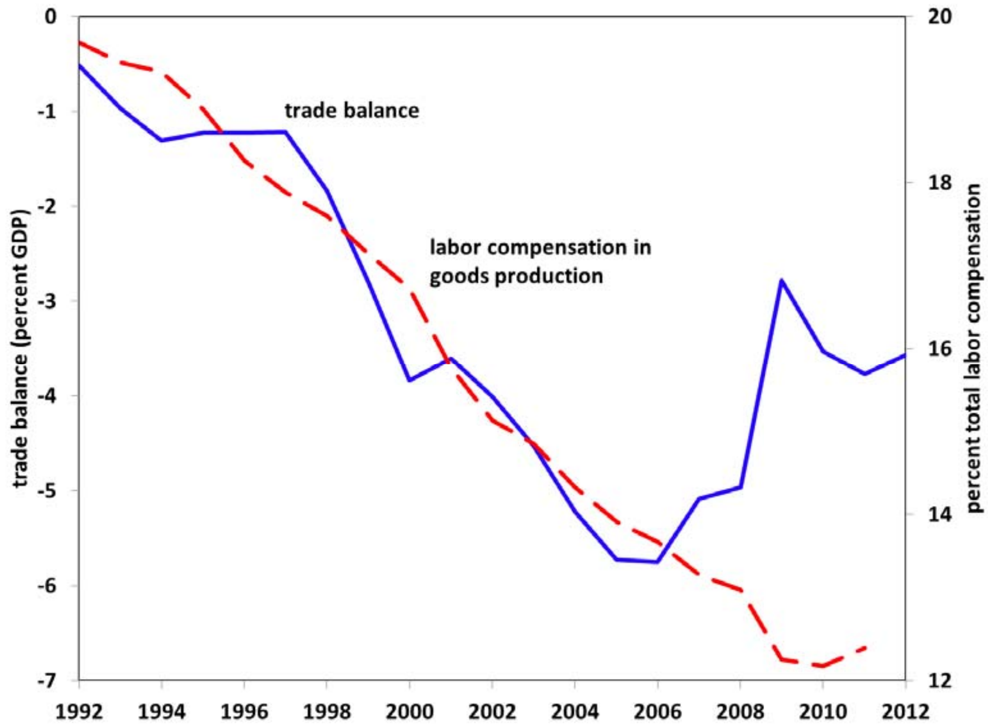

Figure 2: Labor productivity in goods, services, and construction

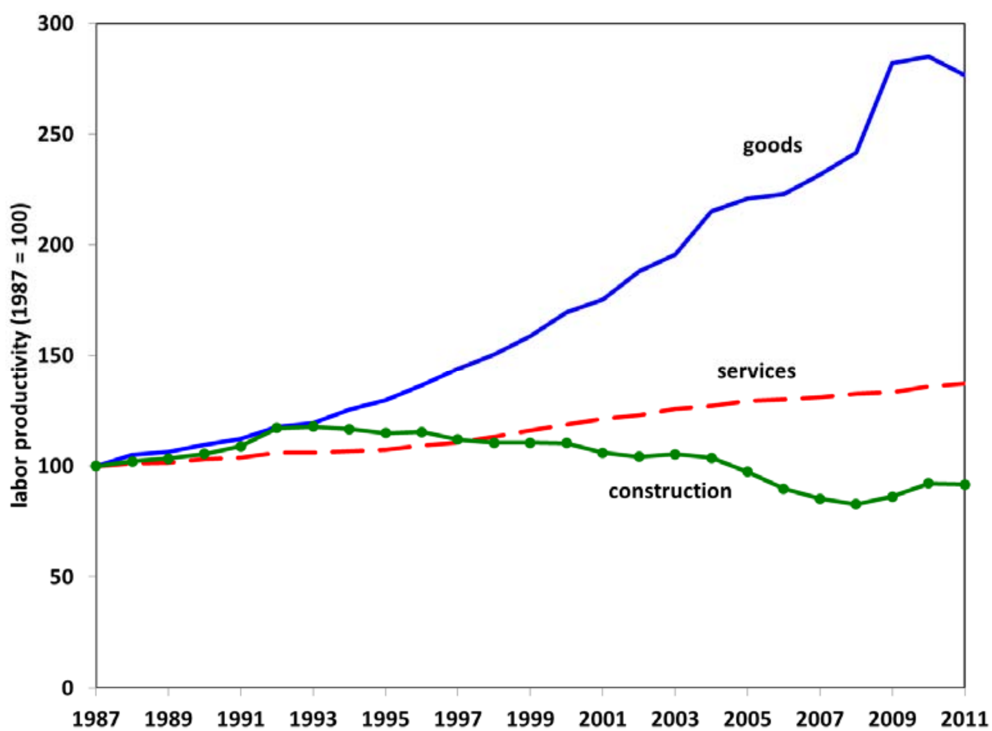

Figure 3: U.S.trade balance, current account balance, and real exchange rate

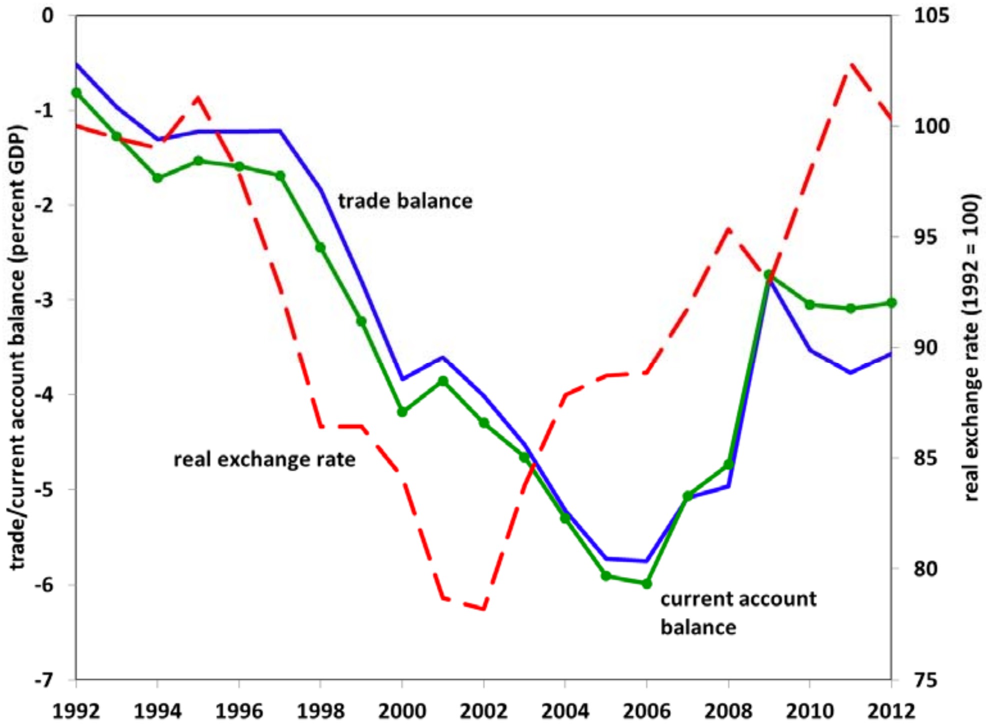

Figure 4: Goods and services trade balances in the model and the data

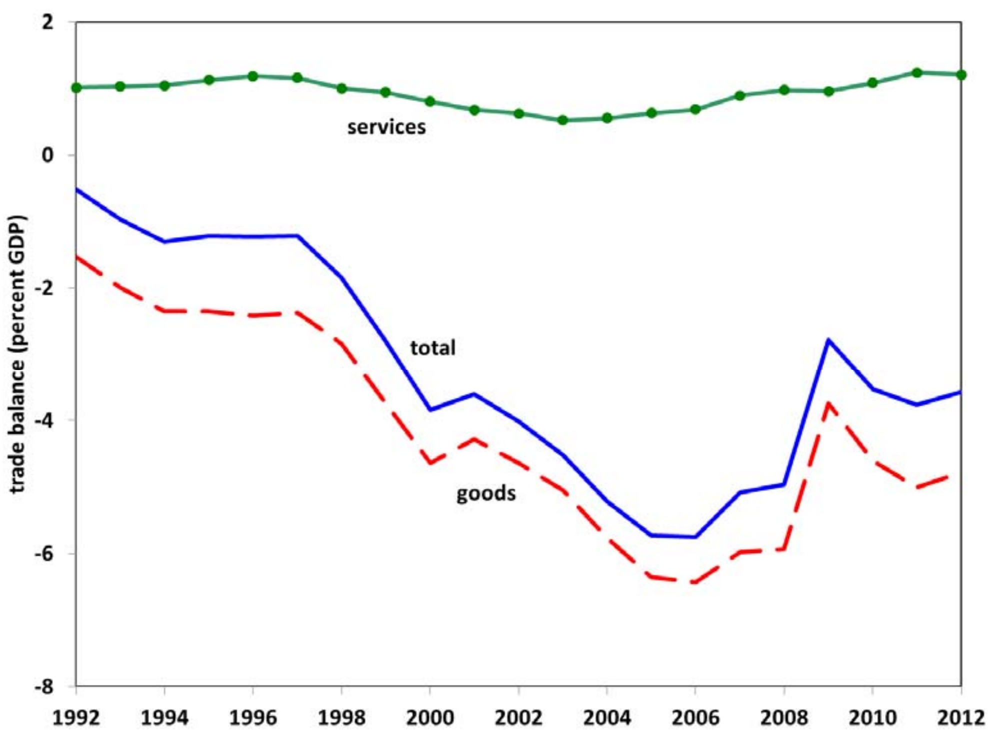

Figure 5: Employment shares in goods and construction sectors

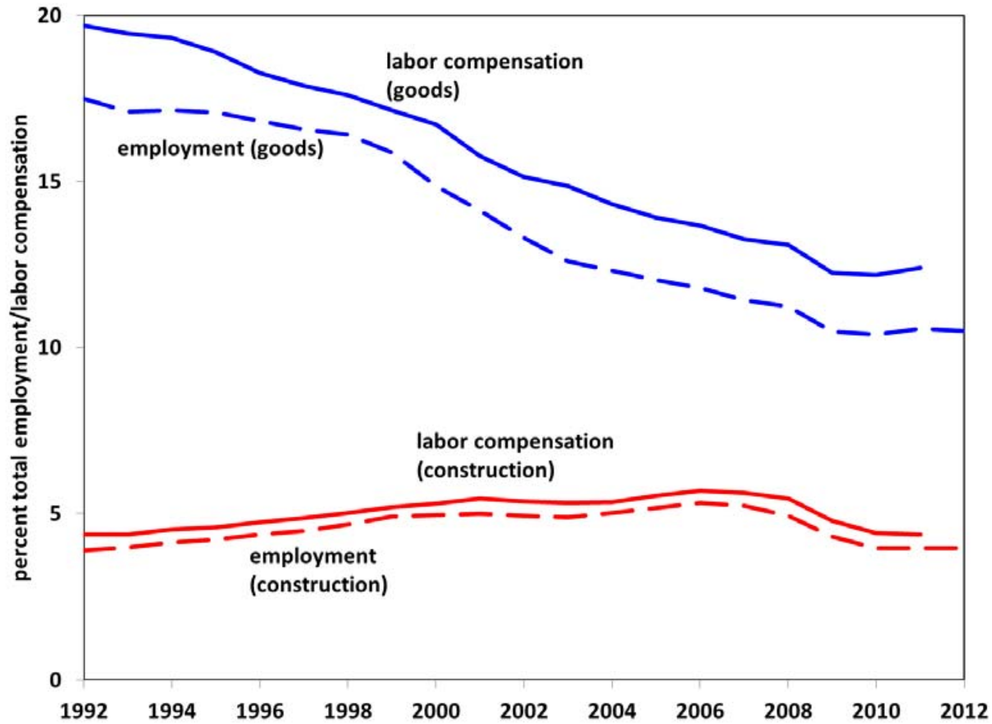

Figure 6: U.S.trade balance with and without the savings glut

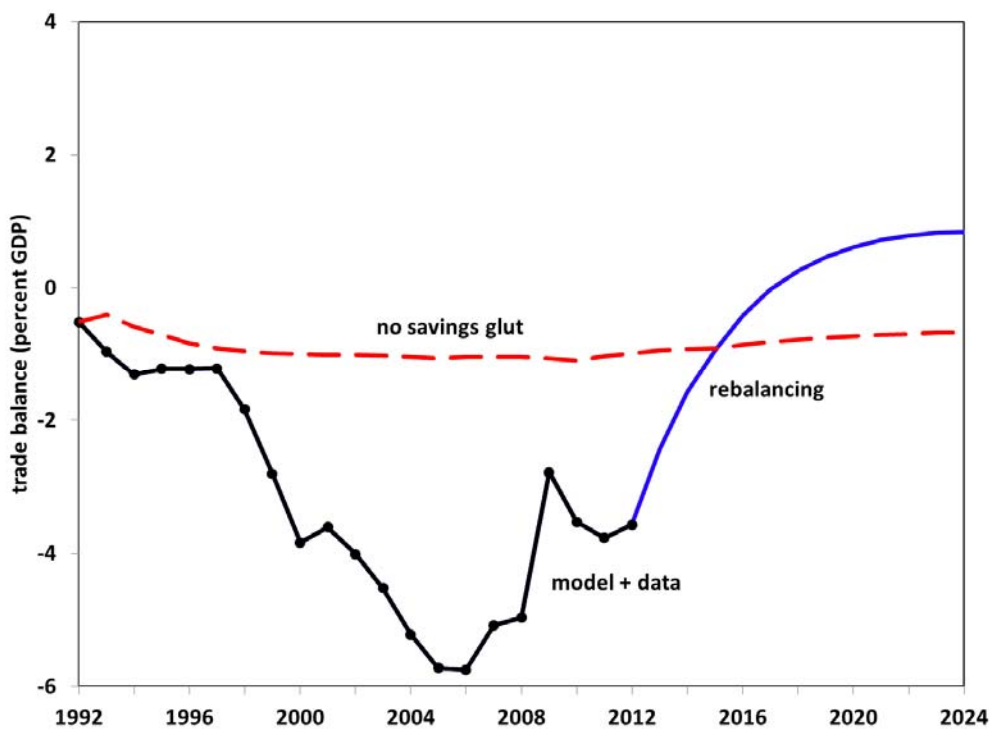

Figure 7: U.S. real exchange rate in the model and the data

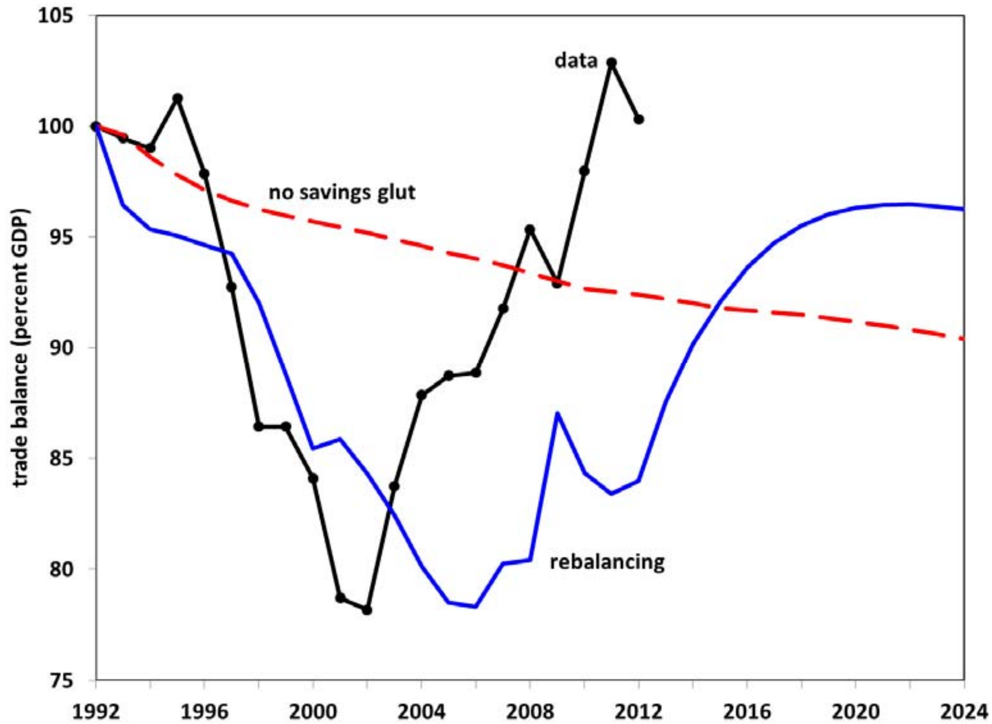

Figure 8: Disaggregated trade balances in the model and the data

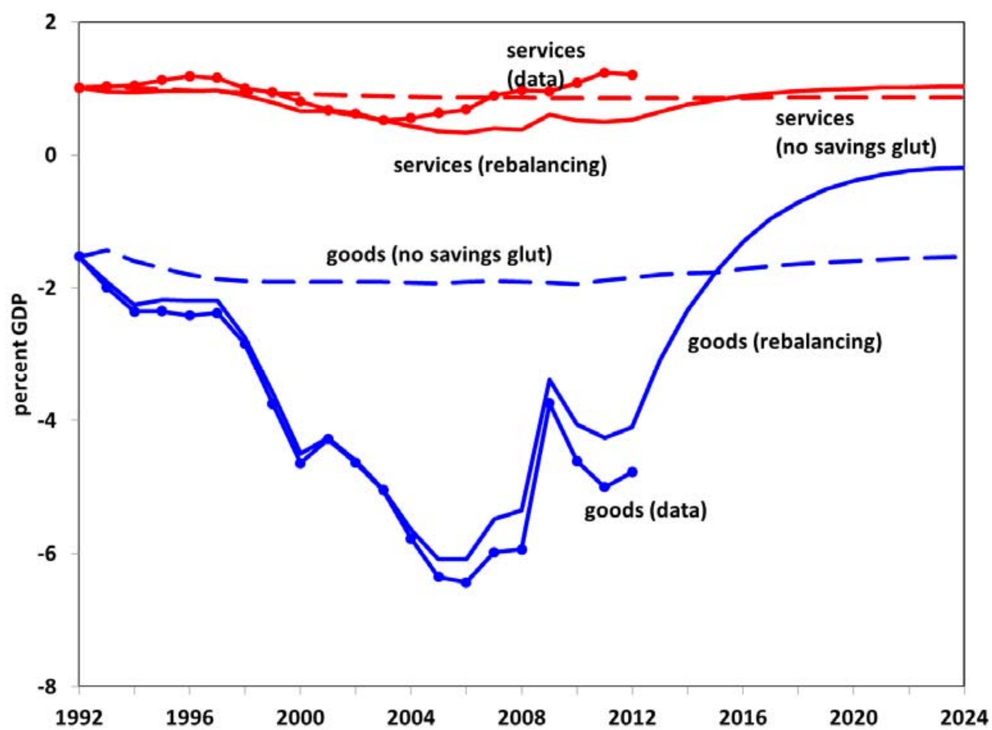

Figure 9: Sectoral employment share in the model and the data

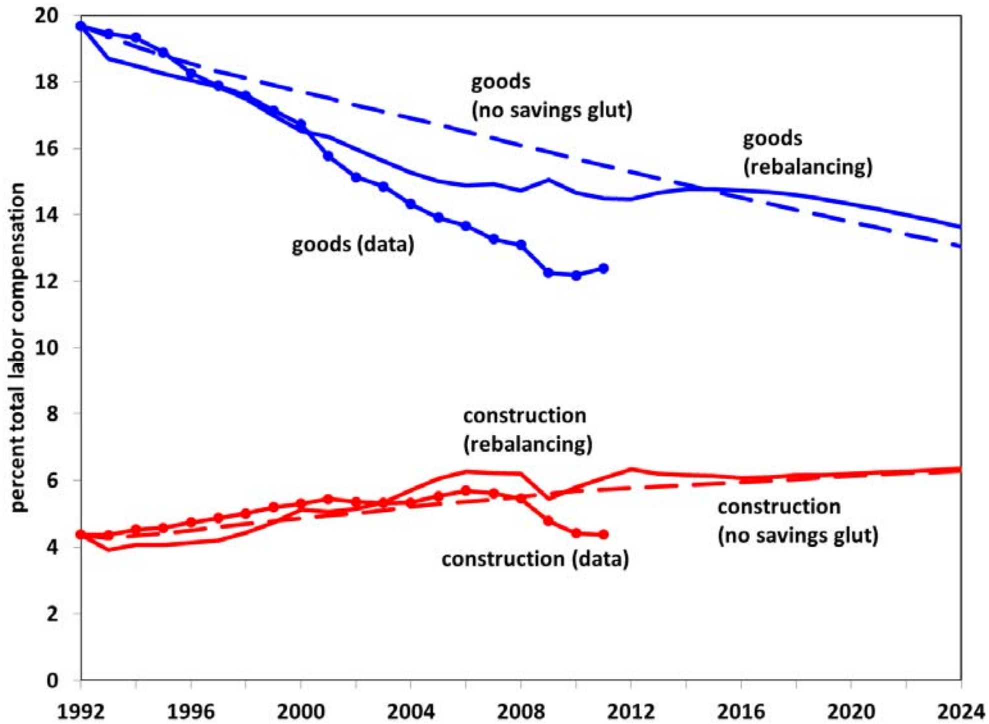

Figure 10: US real exchange rates with China and other trade partners

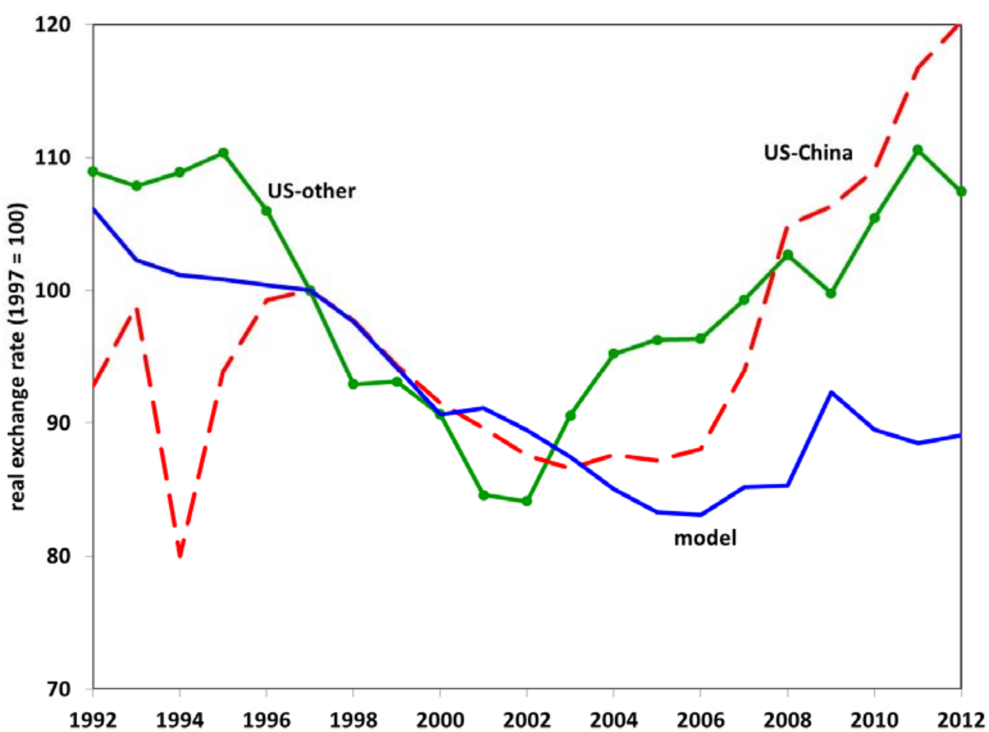

Figure 11: US real interest rates in model and data

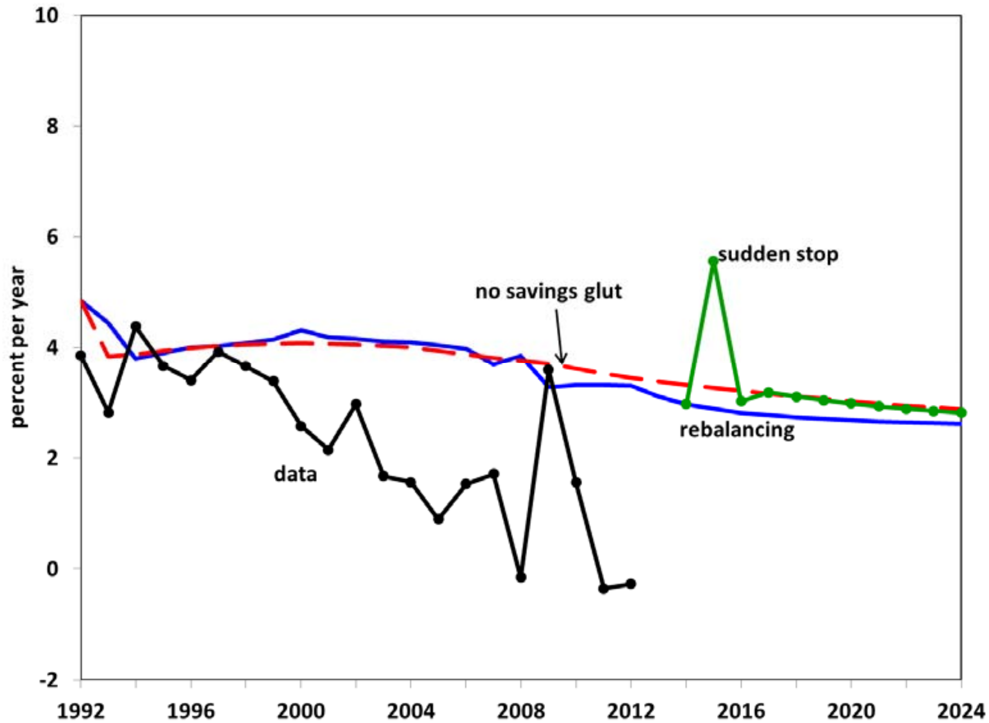

Figure 12: U.S.trade balance and real exchange rate with sudden stop

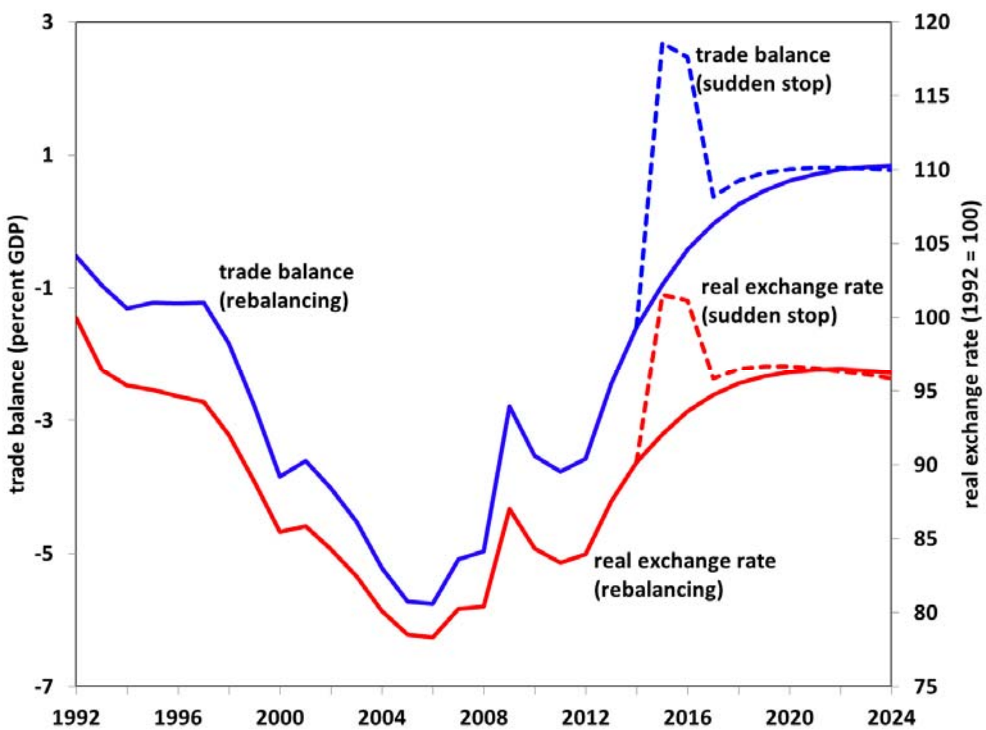

Figure 13: Goods and services trade balances with sudden stop

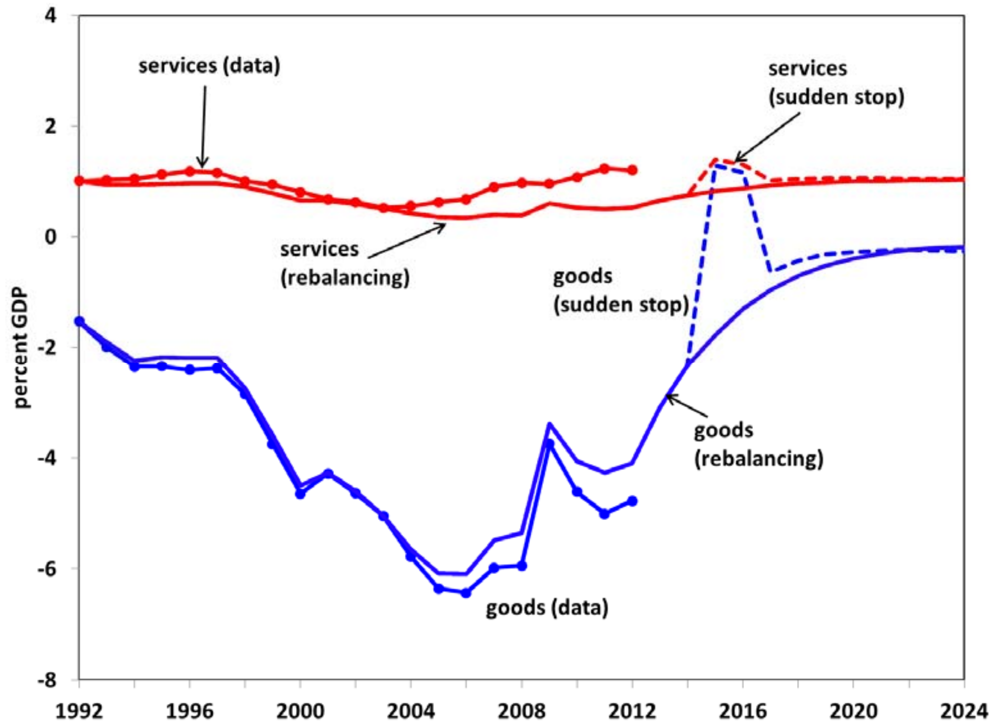

Figure 14: Goods and construction employment shares with sudden stop

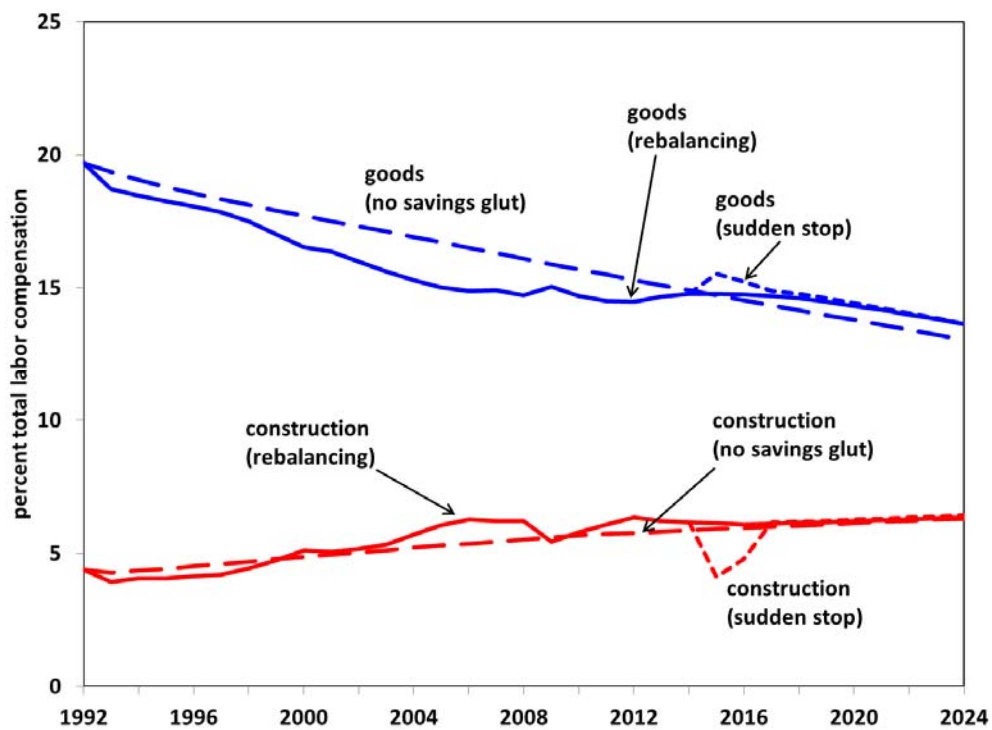

Figure 15: Goods employment with and without intermediate input

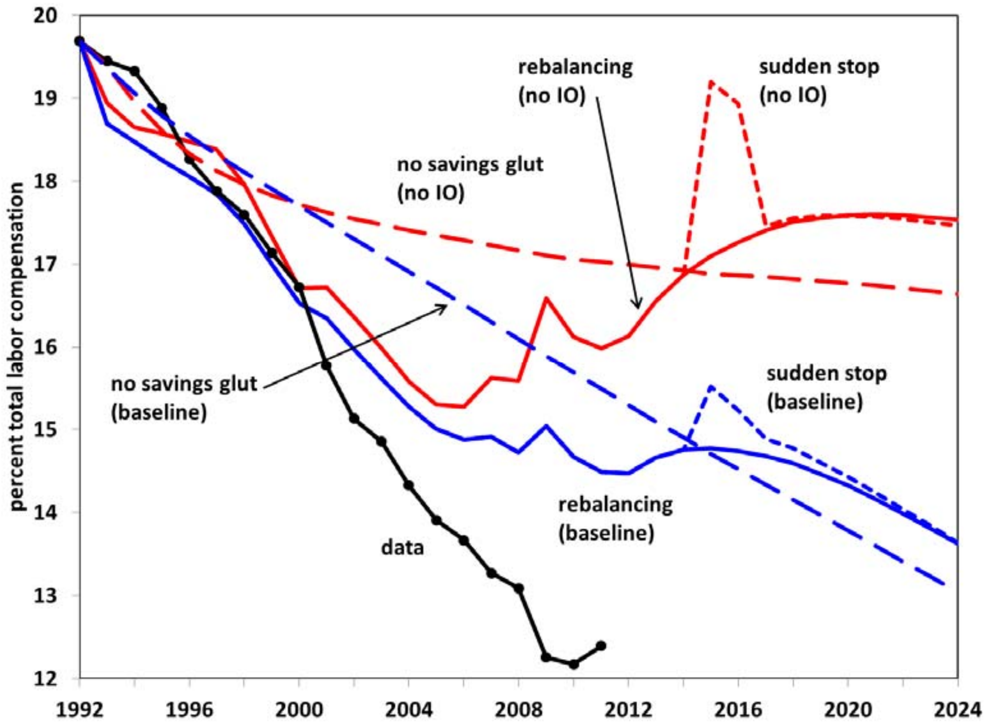

Figure 16: Goods employment with nontradable services

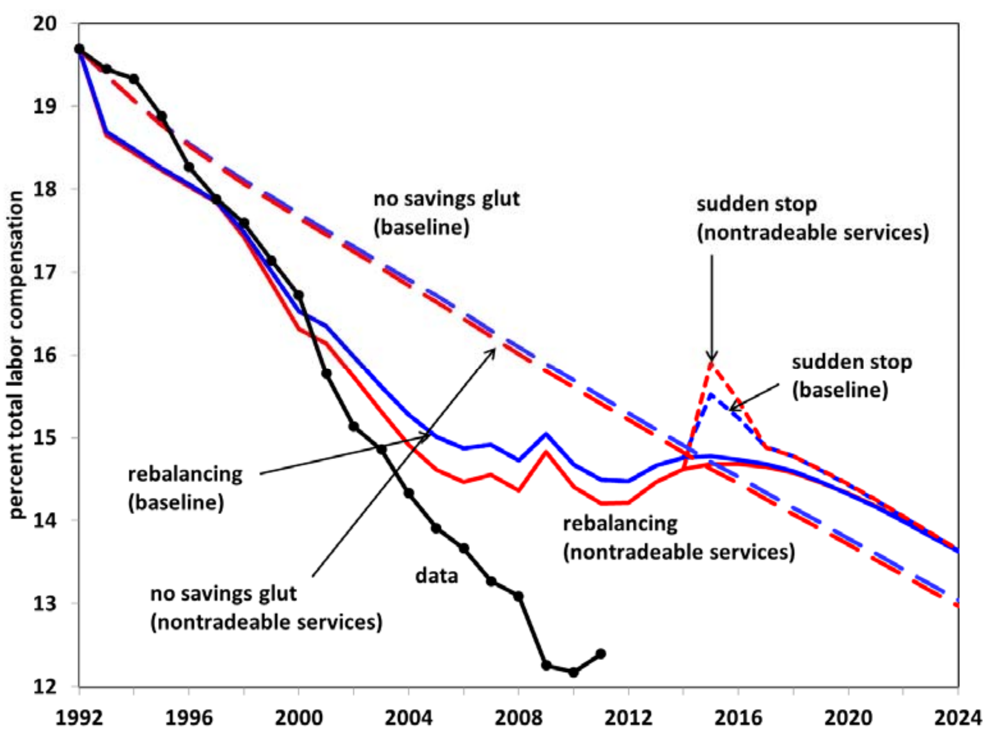

Figure 17: Investment in the baseline and savings drought models

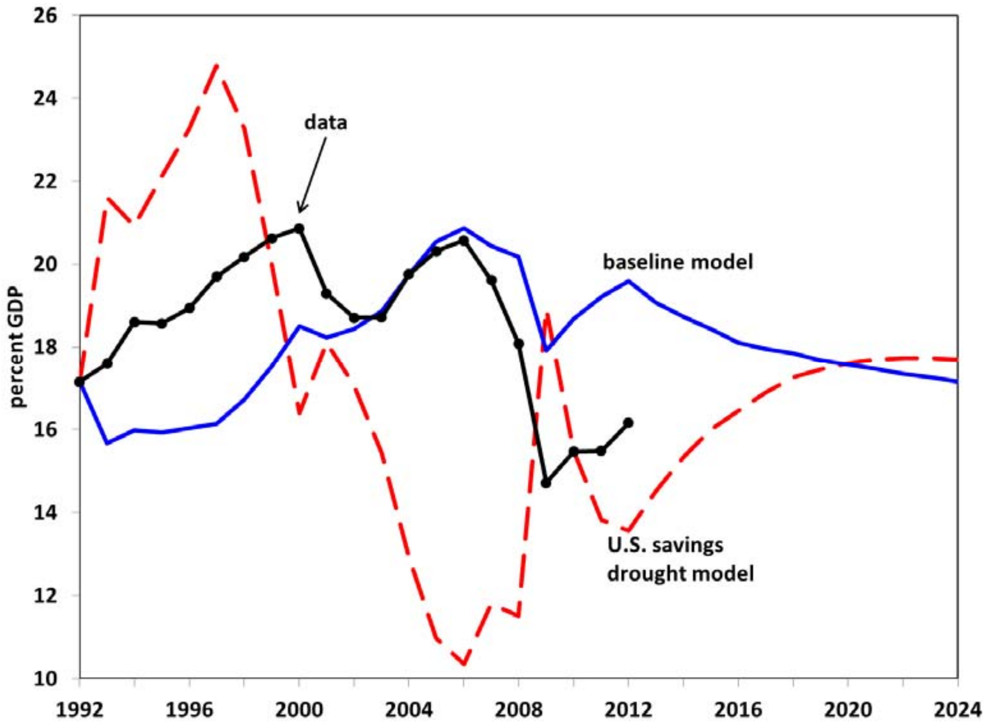

Figure 18: Construction employment in the baseline and savings drought models