

Two Perspectives on Preferences and Structural Transformation*

Berthold Herrendorf

Richard Rogerson

Arizona State University

Arizona State University and NBER

Ákos Valentinyi

National Bank of Hungary, HAS and CEPR

October 3, 2009

Abstract

We ask what specification of preferences can account for the changes in the expenditure shares of broad sectors that are associated with the process of structural transformation in the U.S. since 1947. Following the tradition of the expenditure systems literature, we first calibrate utility function parameters using NIPA data on final consumption expenditure. We find that a Stone–Geary specification fits the data well. While useful, this exercise does not tell the researcher what utility function to use in a model that posits sectoral production functions in value added form. We therefore develop a method to calculate the value added components of consumption categories that are consistent with value added production functions, and use these data to calibrate a utility function over sectoral consumption value added. We find that a Leontief specification fits the data well. Interestingly, the two specifications display very different properties: for final consumption expenditure income effects are the dominant force behind changes in expenditure shares whereas for consumption value added relative price effects are dominant.

Keywords: income effects; preferences; relative price effects; structural transformation.

JEL classification: O11; O14.

*The authors thank Stuart Low and Ed Prescott for many helpful conversations about the topic of this paper. For comments and suggestions, the authors thank Paco Buera, Joe Kaboski, and the conference participants at the SED meetings in Prague (2007) and Istanbul (2009), the conference of the Society for the Advancement of Economic Theory in Kos (2007), the RMM conference in Toronto (2008), the conference on Developments in Macroeconomics at Yonsei University (2009), as well as seminar participants at ASU, Bonn, the San Francisco FED, Southampton, UCLA, Western Ontario, and York (Toronto). For financial support, Herrendorf thanks the Spanish Ministry of Education (Grant SEJ2006-05710/ECON) and Rogerson thanks both the NSF and the Korea Science Foundation (WCU-R33-10005). Hubert Janicki, Loris Rubini, and Paul Schreck provided able research assistance.

1 Introduction

The reallocation of resources across the broad economic sectors agriculture, manufacturing, and services is a prominent feature of economic development. Kuznets (1966) referred to this reallocation as the process of structural transformation and included it as one of the main stylized facts of development. Recent work has argued that modeling structural transformation is important for addressing a variety of substantive issues associated with the evolution of economic aggregates.¹ Despite the extensive work on structural transformation, empirically identifying the key economic forces that shape structural transformation remains an open question. One issue is the role played by preferences. Specifically, while theoretically it is known that households may reallocate expenditure shares across consumption categories because of either income effects or relative price effects, the relative empirical importance of these two forces for structural transformation has not been established.² Resolving this issue has important implications for how policies and technological change influence structural transformation.

In this paper we assess the features of preferences that are quantitatively important in shaping the U.S. structural transformation over the period 1947–2007. Following the tradition of the consumer expenditure systems literature, we begin by using data on final consumption expenditure from the NIPA to calibrate the utility function parameters that provide the best fit to the time series for household consumption expenditure shares on agriculture, manufacturing and services, taking total consumption expenditure and prices as given. We find a strikingly simple result, namely that a Stone–Geary utility function provides a very good fit to these data. Moreover, we show that the nonhomotheticities implicit in this specification are key in shaping the evolution of expenditure shares over the period 1947–2007, implying that income effects rather than relative price effects are the dominant force behind changes in expenditure shares. In terms of the recent theoretical literature on structural transformation, this result supports the

¹See, for example, Laitner (2000) and Gollin et al. (2002) for an application to early development, Messina (2006), Rogerson (2008) and Ngai and Pissarides (2008) for the evolutions of hours worked in Europe and the U.S., Duarte and Restuccia (2009) for productivity evolutions in the OECD, Caselli and Coleman (2001) and Herrendorf et al. (2009) for regional convergence, and Bah (2008) for identifying problem sectors in poor countries. Other contribution to the literature on structural transformation include Echevarria (1997), Kongsamut et al. (2001), Ngai and Pissarides (2007), Acemoglu and Guerrieri (2008), and Foellmi and Zweimuller (2008).

²For example, Kongsamut et al. (2001), assume that only income effects matter, while Baumol (1967) and Ngai and Pissarides (2007), assume that it is only relative price effects.

specification assumed by Kongsamut et al. (2001), but not the one assumed by Baumol (1967) and Ngai and Pissarides (2007).

While at first pass the above result may appear to imply that a Stone–Geary utility function is the appropriate choice for applied studies of structural transformation, we argue that such a conclusion is not warranted because there is an alternative tradition that is also commonly used to define commodities in this context. Specifically, for the exercise reported above, we followed the tradition of the expenditure systems literature and identified the commodities with categories of final consumption expenditure. For example, spending on food from the supermarket will be in agriculture, spending on clothing from a department store will be in manufacturing, and spending on airplane travel will be in services. In contrast, the literature on sectoral productivity studies follows a different tradition. It starts with value added production functions at the sectoral level, and then studies the properties of these functions, including such things as factor input shares and technological change. If technology is specified in this tradition, then consistency dictates that the commodities in the utility function are the value added components of consumption, rather than the final expenditure components. For example, the purchase of a cotton shirt will now be treated as representing consumption from all three categories: cotton from agriculture, processing from manufacturing, and transport and retail from services. We term this specification of commodities the consumption value added specification. If a multi-sector general equilibrium model is specified in this tradition, there is no reason to believe that the utility function calibrated above is appropriate.

We next turn our attention to estimating the utility function when commodities are specified according to consumption value added. It turns out that implementing the estimation is now relatively difficult because data on consumption value added by sector are not readily available. To be sure, data on total value added by sector are readily available, but these data are not sufficient because not all of total value added is consumed. One of the contributions of this paper is to lay out and implement a procedure for extracting the consumption component of total value added and to produce an annual time series for U.S. consumption value added by sector between 1947 and 2007. To the best of our knowledge, we are the first to calibrate a specification of preferences based on these data.

While the household expenditure shares based on consumption value added display the same qualitative properties of structural transformation as emphasized by Kuznets, the calibrated utility function is now strikingly different from the earlier case. In particular, whereas with final consumption expenditure we found that the Stone–Geary specification provided an excellent fit to the data, we now find that nonhomotheticities are relatively unimportant and that a homothetic specification with no substitution between categories (i.e., a Leontief specification) provides a good fit to the data. In other words, contrary to our earlier results, it is now the preference specification adopted by Baumol (1967) and Ngai and Pissarides (2007), rather than the one adopted by Kongsamut et al. (2001), that is consistent with the data. With these preferences, the dominant force behind changes in expenditure shares is relative price effects, rather than income effects.

It is important to emphasize that our results should not be interpreted as saying that researchers can freely rationalize very different choices of utility functions. Our two calibrated utility functions are based not on two different data sets but rather on two different representations of the same underlying data. The final consumption expenditure data are linked to the consumption value added data through complicated input–output relationships, which implicitly translate the income effects that dominate with final consumption expenditure into the relative price effects that dominate with consumption value added, and vice versa. Which specification of the utility function is appropriate depends on the underlying definition of commodities. In the discussion section in the body of the paper, we will provide intuition for the qualitative properties of this translation, and for how it reconciles the seemingly contradictory results of the two methods.

The key point to be taken away is that in the context of multi–sector general equilibrium models, one needs to be careful about how commodities are defined, and in particular about consistency of measurement on the household and production sides of the economy. Moreover, one needs to be careful about importing parameter estimates across models that have different underlying definitions of commodities, even though they may use the same labels. For example, our results show that it is not appropriate to use the utility function that was calibrated from final consumption expenditure together with value added production functions at the sector level.

If one wants to use a utility function that was calibrated from final consumption expenditure, then one either needs to write down a production structure that captures the complexities of the input–output relationships at the sector level, or one needs to find a representation of production that isolates the contribution of capital and labor to the production of final expenditure categories. While this can be done, it is much more difficult than working directly with sectoral value added production functions.³

An outline of the paper follows. In the next section we describe the model and the method that we use to calibrate preference parameters. In Section 3 we describe the final consumption expenditure method and we report the estimation results for this method. In Section 4, we turn to consumption value added. We explain in some detail how to construct the relevant time series of variables from existing data and we report the estimation results. Section 5 compares the results of both methods and provides some intuition for the differences. Moreover, it discusses the relative merits of the two methods and some additional measurement issues. Section 6 concludes. An appendix contains the details about our data work.

2 Model

As noted in the introduction, our objective is to determine what form of preferences for a stand–in household defined over the three broad categories agriculture, manufacturing and services are consistent with U.S. data for expenditure shares over the period since 1947. This section develops the model that we use to answer this question.

We consider an infinitely lived household with preferences represented by a utility function of the form:

$$\sum_{t=0}^{\infty} \beta^t U(c_{at}, c_{mt}, c_{st}, 1 - h_t),$$

where the indices a , m , and s refer to the three broad sectors of agriculture, manufacturing, and services.⁴ Hours of work for the household in period t are denoted by h_t and, with the total time

³Valentinyi and Herrendorf (2008) showed how to construct sectoral production functions that use only capital and labor to produce final expenditure by broad category.

⁴The exact definition of these sectors for each of the two specifications that we consider will be provided later. We note here that we have followed the convention of using the label “manufacturing” to describe a sector which consists of manufacturing and some other sectors (e.g., mining and construction). While the label “industry” is

endowment normalized to 1, the term $1 - h_t$ represents leisure in period t . We will assume that the function U takes one of two forms:

$$U(c_{at}, c_{mt}, c_{st}, 1 - h_t) = u(c_{at}, c_{mt}, c_{st})v(1 - h_t)$$

or

$$U(c_{at}, c_{mt}, c_{st}, 1 - h_t) = u(c_{at}, c_{mt}, c_{st}) + v(1 - h_t).$$

The key feature of these two forms is that time devoted to work has no effect on relative marginal utilities of consumption within a given period, so it will not influence the optimal allocation of expenditures across consumption categories for given prices and total expenditure.

We further assume that the utility function $u(c_{at}, c_{mt}, c_{st})$ is of the form:

$$u(c_{at}, c_{mt}, c_{st}) = \frac{\left[\left(\sum_{i=a,m,s} (\omega_i)^{\frac{1}{\sigma}} (c_{it} + \bar{c}_i)^{\frac{\sigma-1}{\sigma}} \right)^{\frac{\sigma}{\sigma-1}} \right]^{1-\rho} - 1}{1 - \rho},$$

where the ω_i are non-negative weights that add up to one, $\rho > 0$ is the intertemporal elasticity of substitution of consumption, and the $\bar{c}_i$'s are constants. We restrict $\bar{c}_m$ to be zero but allow the $\bar{c}_a$ and $\bar{c}_s$ to take any value. If the $\bar{c}_i$'s are all zero, then preferences are homothetic and $\sigma > 0$ is the within period elasticity of substitution between consumption goods.

A few comments on our specification of the utility function are in order. First, it nests the utility specifications of both Ngai and Pissarides (2007) and Kongsamut et al. (2001). The preferences used by Ngai and Pissarides are the special case in which $\sigma < 1$ and $\bar{c}_a = \bar{c}_s = 0$ whereas the preferences used by Kongsamut et al. are the special case in which $\sigma = 1$, $\bar{c}_a < 0$ and $\bar{c}_s > 0$.⁵

Second, we have followed Kongsamut et al. (2001) and restricted $\bar{c}_m$ to equal zero. We have experimented with an unrestricted specification where $\bar{c}_m$ could take any value, but found that the goodness of fit hardly changed.

sometimes used to describe this sector, we will later use the term “industry” to describe a generic production activity and the index i to denote a generic sector. In view of this, “manufacturing” seemed a better choice.

⁵The preferences used by Kongsamut et al were originally introduced by Stone (1954) and Geary (1950-1951). The implied demand model is often called the Linear Expenditure System. Deaton and Muellbauer (1980) is a subsequent classic contribution to the literature on expenditure systems.

Third, the above specification assumes that σ is the same among all three consumption categories, which may seem somewhat restrictive. Instead, we could have considered more general utility functions with two levels of CES aggregators, each with (possibly) different elasticities of substitution.⁶ We did not pursue this possibility further because we found that our parsimonious class of utility functions includes specifications that provide a very good fit to the expenditure shares.

Lastly, as we show in Appendix A, if all individuals have preferences of the above form and have total consumption expenditure that exceed a minimum level, then aggregate expenditures are consistent with those for a stand-in household with preferences of the same form.⁷ This property extends to settings in which individuals make consumption/savings decisions if there are complete markets.

Consider the stand-in household in a setting in which it maximizes lifetime utility given a market structure that features markets for each of the three consumptions, a labor market, and a market for borrowing and lending at each date t .⁸ Our strategy is to focus solely on the implications for optimal consumption behavior within each period. The advantage of this “partial” approach is that we do not have to take a stand on the exact nature of intertemporal opportunities available to the household (i.e., the appropriate interest rates for borrowing and lending), or to specify how expectations of the future are formed. With these assumptions, if C_t is observed total expenditure on consumption in period t and prices are given by p_{it} , then it follows that the consumption choices in period t must solve the following static optimization problem:

$$\max_{C_{at}, C_{mt}, C_{st}} u(C_{at}, C_{mt}, C_{st}) \quad \text{s.t.} \quad \sum_{i=a,m,s} p_{it} C_{it} = C_t.$$

Assuming interior solutions, the first-order conditions for the above maximization problem are easily derived.⁹ Some simple algebra yields the following expression for the expenditure

⁶See Sato (1975) for a characterization of general CES utility functions.

⁷The precise condition is (1) in Appendix A.

⁸More generally, we could assume some uncertainty in the economy and allow for a richer set of assets that can be traded; see for example Atkeson and Ogaki (1996). What matters for our method is that there are spot markets at each date t for each of the three consumption categories.

⁹In general, of course, the nonhomotheticity terms in our class of utility functions can lead to corner solutions. However, this is not relevant for *aggregate* consumption in a *rich* country such as the postwar U.S. Looking ahead, indeed we will find that the stand-in household chooses quantities that are far away from corners.

shares:

$$s_{it} \equiv \frac{p_{it}c_{it}}{C_t} = \frac{1 + \sum_{j=a,m,s} \frac{p_{jt}\bar{c}_j}{C_t}}{\sum_{j=a,m,s} \frac{\omega_j}{\omega_i} \left(\frac{p_{jt}}{p_{it}} \right)^{1-\sigma}} - \frac{p_{it}\bar{c}_i}{C_t}.$$

In the empirical worked reported below, we will calibrate the parameters of the utility function by minimizing the distance between the expenditure shares in the data and those generated by the model. While we have also calibrated the parameters of the utility function via maximum likelihood, we nonetheless will only report the results from the minimum-distance exercise. The reason for this is that, as we will see below, our data series are nonstationary, and the statistical properties of the maximum likelihood estimator are unknown in this case.¹⁰ In any case, we note that the parameter values generated from the minimum distance exercise and the maximum likelihood estimation are very similar in almost all cases.

3 Final Consumption Expenditure

The final consumption expenditure method originated in the literature on expenditure systems. It associates the arguments of the utility function with final expenditure of households over different categories of goods and services. Specifically, this method classifies the expenditures on individual *commodities* into the three broad sectors of agriculture, manufacturing, and services. For example, purchases of food from supermarkets will be included in c_{at} , purchases of clothing from retail establishments will be included in c_{mt} , and purchases of air-travel services will be included in c_{st} .

3.1 Implementing the Final Consumption Expenditure Specification

The required data in this case are total consumption expenditure and the expenditure shares and prices for final consumption expenditure on different commodities. The exact data sources can be found in Appendix B.1. While expenditure shares do not depend on how one splits total expenditures into their price and quantity components, the series for prices do. That is, given total

¹⁰For further discussion on these issues, see the review of Barnett and Serletis (2008).

expenditure, different procedures for inferring the consumption quantities will imply different relative prices. Consistent with BEA measurement, we measure final consumption quantities using chain-weighted methods. For the period 1947–2007 and for the available commodities, we obtain annual data on final consumption expenditure, chain-weighted final consumption quantities, and prices from the BEA. Since quantities calculated with the chain-weighted method are not additive, we use the so called cyclical expansion procedure to aggregate quantities that are not available from the BEA.¹¹ We assign each commodity to one of the three broad sectors agriculture, manufacturing, and services. A detailed description of this assignment can be found in Appendix B.2. Note that for estimating utility function parameters we do not need to know whether the commodities purchased by the household are produced in the U.S. economy or imported. All that matters for our exercise is information on total consumption expenditure, expenditure shares and prices.

Figures 1–3 show the resulting evolution of the expenditure shares, prices and quantities, respectively. A few remarks regarding the figures are in order. Looking at Figure 1, we see that the data are consistent with the standard (asymptotic) pattern of structural transformation: The expenditure share for services is increasing, while those for agriculture and manufacturing are decreasing. Turning next to Figure 2, which shows the evolution of prices (with prices in 1947 normalized to 1), we see that while all three prices have increased, the price of services has increased relative to both manufacturing and agriculture and the price of agriculture has increased relative to manufacturing. Figure 3 shows real quantities relative to their 1947 values. Here we see that while the quantities of all three categories have increased, the quantity of manufacturing has grown the most, while the quantity of agriculture has grown the least.

Figures 1–3 already suggest some of the qualitative features of the utility specification that our estimation will select. First, note that the price of services has increased relative to that of agriculture, while at the same time the quantity of services has also increased substantially relative to that of agriculture. This is qualitatively inconsistent with a homothetic utility specification, which would have relative prices and relative quantities move in opposite directions. In the context of our class of utility functions, reconciling these observations amounts to having

¹¹See Landefeld and Parker (1997) for more details.

Final Consumption Expenditure Per Capita

Figure 1: Expenditure shares




Figure 2: Price Indices (1947=1)




Figure 3: Quantity Indices (2000 chained dollars, 1947=1)




Table 1: Results with Final Consumption Expenditure

	(1)	(2)	(3)
σ	.81	1	.20
$\bar{c}_a$	-1,208	-1,182	0
$\bar{c}_s$	8,024	15,999	0
ω_m	.18	.15	.19
ω_s	.81	.83	.70
Loss	.105	.112	.342

$\bar{c}_a < 0$ and/or $\bar{c}_s > 0$. Second, as the price of agriculture relative to manufacturing has increased, the quantity of agriculture relative to manufacturing has decreased. This is consistent with there being substitutability between agriculture and manufacturing. While to some extent this could also be accounted for by having $\bar{c}_a < 0$, in the context of our preference specification, it turns out that σ will come out close to one.

3.2 Results with Final Consumption Expenditure

We now calibrate the parameters of the utility function by minimizing the unweighted squared differences between the actual and the predicted expenditure shares, given the data for prices and total final consumption expenditure.

Table 1 shows the results for three different specifications. In particular, Column (1) shows the results when we do not impose any restrictions on the parameters other than $\bar{c}_m = 0$. As can be seen, the implied value of σ is 0.81. Moreover, the signs of the two unrestricted non-homothetic terms have the pattern that Kongsamut et al. (2001) suggested, that is, $\bar{c}_a < 0$ and $\bar{c}_s > 0$. The nonhomotheticities imply that holding relative prices fixed, as total expenditure grows over time, the consumption share of agriculture will go down while that of services will go up.¹² Figure 4 shows that the fit of the calibrated model from Column (1) to the data on final consumption expenditure shares is very good indeed.

While the specification from Column (1) is similar to the Stone–Geary specification imposed by Kongsamut et al. (2001), it is not identical, since they assumed $\sigma = 1$. It is therefore

¹²This result is broadly consistent with what panel studies of household consumption find; see Houthakker and Taylor (1970) for a classic contribution.

Fit of Expenditure Shares with Final Consumption Expenditure

Figure 4: Fit of Column (1)





Figure 5: Fit of Column (2)



interesting to assess the extent to which the specification of Kongsamut et al. (2001) fits the data. To examine this, the second column redoes the loss minimization procedure, this time imposing that $\sigma = 1$. The nonhomothetic terms retain the same sign configuration, although the magnitude of $\bar{c}_s$ changes quite substantially. This is intuitive because with a higher σ households react to the large increase in the relative price of services by substituting away from services. The higher $\bar{c}_s$ term compensates for this. Figure 5 shows that the specification of Column (2) fits virtually as well as the specification of Column (1). This is consistent with the fact that the loss reported in Table 1 hardly increases.

Several papers estimated the linear expenditure systems that are implied by the Stone–Geary utility specification. Our results are most closely related to those of a literature that used time series data for final consumption expenditure per capita. A prominent example is Pollak and Wales (1969), who used US data 1948–1965 on the four broad categories food, clothing, shelter, and miscellaneous. Pollak and Wales also found that the linear expenditure system fits the data very well and that nonhomotheticity terms are important.¹³

We conclude that when using data on final consumption expenditure, the data broadly support the Stone–Geary specification of Kongsamut et al. (2001). Having said that, note that they

¹³For a subsequent literature review, see Blundell (1988).

Table 2: Nonhomotheticity Terms Relative to Final Consumption Expenditure from the Data

	1947	2007
$-p_a \bar{c}_a / C$	.17	.04
$p_s \bar{c}_s / C$	1.22	.51
$-\bar{c}_a / c_a$	.80	.57
$\bar{c}_s / c_s$	2.51	.71

also imposed the condition

$$p_{at} \bar{c}_a + p_{st} \bar{c}_s = 0,$$

which is required for the existence of a generalized balanced growth path in their model.¹⁴ This condition is rather trivially not consistent with the data we use. The simplest way to see this is to look at Figure 2, which clearly shows that p_{st}/p_{at} has been steadily increasing since 1947, whereas, of course, $\bar{c}_a$ and $\bar{c}_s$ are constants.

3.3 Income versus Price Effects with Final Consumption Expenditure

It is of interest to look more closely at the importance of income versus relative price effects in accounting for the observed changes in the shares of final consumption expenditures. To put the size of the calibrated nonhomotheticity terms of Column (1) into perspective, Table 2 reports the $\bar{c}_i$'s relative to final consumption expenditure from the data. Most importantly, rows three and four show that in 1947 both nonhomotheticity terms were sizeable compared to the actual consumption quantities of agriculture and services. This suggests that income effects play a very important role in shaping the shares of final consumption expenditure.

To explore further the importance of income versus relative price effects, Figure 6 shows the fit of the expenditure shares implied by the parameters of Column (1) when total expenditure changes as dictated by the data but relative prices are held constant at their 1947 values. We can see that although the fit deteriorates somewhat, the model still captures the vast majority of

¹⁴Given the nonhomotheticity terms, their model does not have a balanced growth path in the usual sense of the word. They therefore consider a generalized balanced growth path, which they define as a growth path along which aggregate variables such as the real interest rate are constant.

Income Effects with Final Consumption Expenditure Shares

Figure 6: Fit of Column (1) with Relative Prices Fixed at 1947 Values





Figure 7: Fit of Homothetic Specification in Column (3)



the changes in the expenditure shares. The main discrepancies between the data and the model are that the share of services now increases somewhat more than in the data and the share of agriculture now decreases somewhat more than in the data. These discrepancies are intuitive since the price of services increases relative to agriculture during our sample period, and would therefore work to partially offset the changes associated with the income effects.

A second way of judging the importance of income versus relative price effects is to ask how well a homothetic specification can fit the data, since such a specification necessarily implies that total expenditure has no effect on expenditure shares. Column (3) presents the results of a specification in which the nonhomothetic terms are restricted to equal zero. We can see that the calibrated elasticity term σ drops from 0.81 to 0.21. Moreover, the loss more than triples, suggesting that the fit gets considerably worse. Figure 7 confirms this by showing that the fit becomes fairly bad for agriculture relative to the previous two specifications. Absent income effects and given a rising relative price of services, the way in which the homothetic specification can account for the large rise in the expenditure share for services is by having very little substitutability. But, as noted earlier, absent income effects, fitting the expenditure share for agriculture requires some substitutability. Hence, the model without income effects cannot do a good job of matching all expenditure shares.

We conclude that the income effects associated with the nonhomotheticities are the dominant source of the observed structural transformation in the shares of final consumption expenditure.

4 Consumption Value Added

Many multi-sector general equilibrium models represent the sectoral production functions in value added form, in which case the arguments of the utility function necessarily represent the value added components of final expenditure. Individual *industries* are then classified into the three broad sectors agriculture, manufacturing, and services, and a sector is a collection of industries, with sector value added being the sum of the value added of the industries belonging to it. Effectively, this way of proceeding breaks consumption spending into its value added components. For example, purchases from supermarkets will then be broken down into the components of c_{at} (food), c_{mt} (processing of the food) and c_{st} (retail services of the supermarket). Similarly, purchases of clothing will be broken down into the components of c_{at} (raw materials, say cotton), c_{mt} (processing of cotton into clothing) and c_{st} (retail services), and purchases of air-travel services will be broken down into the components of c_{mt} (fuel) and c_{st} (transportation services).

Whether one prefers to use final consumption expenditure or consumption value added will depend on data availability and the specific application. While we will discuss the relative merits of each method in more detail in subsection 5.2 below, there are two key points that we want to emphasize at this point. First, there is no reason to believe that the parameters of the utility function are invariant to the definition of commodities. Looking ahead, the distinction between the two different specifications will turn out to be very significant, since we will find that they imply preferences with very different qualitative properties. Second, and related, it is important to emphasize that the two specifications are two different representations of the same underlying data. Put differently, the data on final consumption expenditure are linked to the data on consumption value added through complicated input-output relationships, and vice versa. Our results should therefore not be interpreted as implying that different data sets provide

different parameter estimates. Rather, they tell us that different transformations of given data have different properties.

4.1 Implementing the Consumption Value Added Specification

We now describe how to construct the relevant data when one identifies the three consumption categories with their respective value added components. The exact data sources can be found in Appendix B.1. Similar to the case of final expenditure shares, there is annual data available from the BEA on value added by industry, as well as real value added and prices. As we mentioned above, the consumption value added method assigns industries, instead of commodities, to the three broad sectors. Appendix B.2 describes the details of this assignment.

Although readily available, the data on value added and prices are not sufficient for our purposes. The reason is that value added data come from the production side of the national income and products accounts, and so they contain both consumption and investment. It is therefore necessary to devise a method to extract the consumption component from the production value added of each sector. This has not been sufficiently appreciated in the literature, which often proceeds by assuming that all investment is done in manufacturing.¹⁵ While the problem with this assumption may not have been readily apparent in the past, recent data make it clear that this assumption is problematic. The reason for this is that in the BEA data, manufacturing production value added has been consistently smaller than investment since 1999. We therefore need to properly extract the consumption component from the production value added in each sector. One contribution of our paper is to lay out a procedure that achieves this.

To carry out this extraction one needs to combine the value added data from the income side of the NIPA with the final expenditure data from the expenditure side of the NIPA. The complete details of this procedure are fairly involved, and so we relegate its description to Appendix C. Here we provide just a rough sketch of the procedure. A key difference between value added data from the income side and final expenditure data from the expenditure side is that the former are measured in what the BEA calls “producer’s prices” whereas the latter are measured in “purchaser’s prices”. From a practical perspective, the key difference is that

¹⁵Examples include Huffman and Wynne (1999) and Buera and Kaboski (2009).

Consumption Value Added Per Capita

Figure 8: Expenditure Shares




Figure 9: Price Indices (1947=1)





Figure 10: Quantity Indices (2000 chained dollars, 1947=1)



purchaser's prices include distribution costs whereas producer's price do not (distribution costs are sales taxes and transport, wholesale, and retail services). For example, in the case of a shirt purchased from a retail outlet, the purchaser's price is the price paid by the consumer in the retail outlet whereas the producer's price is the price of the shirt when it leaves the factory.

In order to break final consumption expenditure into its value added components the first step therefore is to convert final consumption expenditure measured in purchaser's prices into those measured in producer's prices. This amounts to removing distribution costs from final consumption expenditure on goods and moving them into expenditure on services. Appendix C.1 explains the details of this calculation. Once this is done the second step is to use the input–output tables to determine the sectoral inputs in terms of value added that are required to deliver the final consumption expenditure. This involves an object called the total requirement matrix which is derived from the input–output tables. Appendix C.2 explains the details of this procedure.

Two points are worth stressing. First, since we are interested in the time series properties of consumption value added, and the structure of input–output relationships changes over time, an important feature of our calculation is that we use all annual input–output tables together all benchmark tables that are available for the period 1947–2007. Second, when we break final consumption expenditure into its value added components we follow the BEA and treat imported goods as if they were produced domestically with the same technology that the U.S. uses to produce them. Given this assumption, we again do not have to stand on whether intermediate goods are produced domestically or imported.¹⁶

Having broken final consumption expenditure into its value added components, we obtain data on consumption value added expenditure shares and chain-weighted prices and quantities, which are displayed in Figures 8–10. Note that these figures display the same qualitative pattern for consumption value added shares that we saw in the analogous figure for final consumption expenditure shares. Hence, both representations are consistent with the stylized facts about structural transformation. However, although the shares display similar qualitative behavior, there are some important differences in the behavior of relative prices and quantities. First, Fig-

¹⁶Appendix C.2 explains this point in more detail.

Table 3: Results with Consumption Value Added

	(1)	(2)
σ	0	0
$\bar{c}_a$	-136.7	0
$\bar{c}_s$	3,652	0
ω_m	.15	.19
ω_s	.84	.80
Loss	.12	.25

ure 9 shows that while the price of services still increased the most, the price of manufacturing now increased by more than that of agriculture. Second, Figure 10 shows that relative quantities behave very differently. Whereas Figure 3 indicated substantial changes in relative quantities, Figure 10 suggests that the relative quantities of manufacturing and services now hardly change over the entire period, while the relative quantity of agriculture remains fairly constant after about 1970.

Given that relative prices changed substantially, the near constancy of relative quantities, particularly of manufacturing relative to services, suggests a very low degree of substitutability between the different components of consumption value added. Moreover, the near constancy of the relative agricultural quantity after 1970 suggests that nonhomotheticities will not play as important a role as before. We will return to the significance of these observations below when we present the calibration results.

4.2 Results with Consumption Value Added

We again choose parameters so as to minimize a loss function which is the sum of the squared difference between actual and predicted expenditure shares. Results are contained in Table 3. Column (1) reports the parameter estimates when we impose no restrictions. Strikingly, the value of σ is calibrated to be equal to 0, which in the absence of nonhomotheticities is the Leontief specification. The nonhomothetic terms have the same signs as before. Figure 11 shows that based on the estimates in Column (1), the fit of the model to the expenditure share data is again very good.

Table 4: Nonhomotheticity Terms Relative to Consumption Value Added from the Data

	1947	2007
$-p_a \bar{c}_a / C$	.08	.0047
$p_s \bar{c}_s / C$	.33	.11
$-\bar{c}_a / c_a$	.97	.38
$\bar{c}_s / c_s$	.52	.13

As before, it is of interest to ask how important income and relative price effects are in accounting for the observed changes in the shares of consumption value added. One simple and revealing comparison about the relative importance of income effects for the two different data sets is to look at the size of the calibrated $\bar{c}_i$'s relative to total consumption expenditure from the data. The first two rows of Table 4 show that these ratios are now considerably smaller than in the case of final consumption expenditure, suggesting that income effects will be much less important. However, one needs to be somewhat careful with drawing immediate conclusions because the third row of the table shows that in 1947 the agricultural consumption value added from the data is now fairly close to $\bar{c}_a$. Intuitively, this can be understood by going back to Figure 10 above, which showed that the quantity of agricultural consumption value did not increase in the first part of our sample. Given that $\sigma = 0$, our calibrated utility function delivers this by having agricultural consumption value added close to $\bar{c}_a$ initially.

To explore more systematically the importance of income versus relative price effects, Figure 12 shows the fit of the expenditure shares implied by the parameters of Column (1) when relative prices change as dictated by the data while total consumption expenditures are held constant at their 1947 value. While the fit deteriorates somewhat, it remains reasonably good. This suggests that relative price effects are the dominant force behind the changes in the expenditure shares of consumption value added. A second way of establishing this is to evaluate the ability of a homothetic specification to fit the data. To examine this, Column (2) in Table 3 presents the results under the restriction $\bar{c}_a = \bar{c}_m = \bar{c}_s = 0$. While the loss doubles the relative weights hardly change compared to Column (1). Figure 13 plots the expenditure share predicted by the homothetic specification. We can see that compared to the nonhomothetic specification of Figure 11, the fit remains very good.

We conclude that the consumption value added data broadly support the homothetic preference specification used by Ngai and Pissarides, though in the somewhat extreme form of a Leontief specification, i.e. $\sigma = 0$. Since introspection would suggest substantial willingness to substitute across many commodities, some readers might question the empirical plausibility of preferences that do not permit any substitution across the consumption value added categories agriculture, manufacturing, and services. It is therefore important to understand exactly what the result $\sigma = 0$ means. Although having $\sigma = 0$ implies that there is no substitutability across the three categories, it is completely consistent with there being substantial substitution *within* each of these categories. In particular, since the categories are quite broad, having $\sigma = 0$ does not in any sense imply that there is no substitutability between all the different goods and services that individuals consume.

A simple example may be useful. Most readers will agree that there is some substitutability between the two activities of going to the movies and going to sporting events. When we represent these activities in consumption value added terms, we see that both of them involve some consumption of goods (e.g., the use of buildings) and some consumption of services (e.g., actors and athletes producing entertainment services). To us it seems reasonable to think that the key dimensions of substitution are within these two value added categories, i.e., that the key substitution is between the uses of buildings or the uses of athletes' and entertainers' time, rather than between goods and services per se. While this is not to suggest that one cannot think of specific examples with some substitution between specific goods and specific services, the key point we want to make is that there is likely to be considerably more substitutability within each of the value added categories.

Our conclusion that a low σ provides the best fit to value added data is related to independent research by Buera and Kaboski (2009). These authors ask whether there are parameters for which a canonical model of structural transformation can match the value added shares by sector between 1870 and 2000. They take as given initial GDP, a time series of overall TFP, and proxies for the price indices of the three sectors. They do not use information on investment, total consumption, and output from data, but let their model choose these quantities under the assumption that all model investment is done in the model sector manufacturing. Their

Fit of Expenditure Shares with Consumption Value Added

Figure 11: Fit of Column (1)




Figure 12: Fit of Column (1) with Income Fixed at 1947 Value





Figure 13: Fit of Homothetic Specification in Column (2)



preferred choice for σ comes out fairly low as well ($\sigma = 0.5$), but in contrast to us they conclude that their model cannot provide a good fit to the data.

This raises the question why the conclusions of the two studies are different. Upon closer inspection, there are several potentially important differences. To begin with, Buera–Kaboski consider data at ten-year intervals from 1870 to 2000, which cover a much larger range of expenditure shares than our data. Moreover, since our data are not available prior 1947, they use data for sector expenditure shares and prices that are not necessarily mutually consistent.¹⁷ Lastly, Buera–Kaboski do not calibrate the utility function given the expenditure shares of consumption value added, prices, and total consumption, but they determine it together with the sector production functions so as to match the shares of production value added.

To best evaluate what drives the differences between the conclusions of the two studies, it would be natural to redo our calibration exercise for 1947–2007 using the data of Buera–Kaboski. Unfortunately, this is not possible, since as noted earlier, investment has exceeded manufacturing value added since 1999, implying that we cannot extract the consumption part of production value added following their assumption that all investment is done in manufacturing. Independently of this, we do note that the series used by Buera–Kaboski to proxy for relative prices do display very different behavior than the true relative price series for value added over the period 1947–2007, suggesting that data differences probably play a significant role.

5 Discussion

5.1 Comparing the Results

Both calibration exercises yield utility specifications that provide very good fits to the data. Somewhat surprisingly, in one case it is the Stone–Geary specification used by Kongsamut et al. (2001) that provides a very good fit to the data, whereas in the other case it is the homothetic specification used by Ngai and Pissarides (2007). The importance of relative price and income effects in accounting for changes in expenditure shares are therefore dramatically different in the

¹⁷Buera–Kaboski use the implicit deflator of services in NIPA and the producer price index of finished goods from the BLS. The former is based on gross sales while the latter is based on final expenditure. In contrast, we use price indices that are based on value added.

two cases. In the case of final consumption expenditure, income effects are the dominant force behind changes in the expenditure shares, whereas in the case of consumption value added, relative price effects are the dominant force.

Given the very different results for the calibrated utility specifications in the two cases, it is of interest to ask if there is any intuition that would support the nature of the differences. In particular, we would like to understand why the final consumption expenditure specification exhibits a greater degree of substitution and a more important role for nonhomotheticities. The intuition is sharpest if we focus on two consumption items: food from supermarkets and meals from restaurants, though it easily extends to a larger set of goods and services.

The intuition for greater substitutability in the final consumption expenditure specification comes from the fact that this specification may place items with similar underlying characteristics into different categories. To stay with our example, this method counts food from supermarkets in agriculture, while meals from restaurants are counted in services. One would expect there to be substitutability between the two items because they both use the intermediate input food. In contrast, in the consumption value added specification, all agricultural inputs into food production are counted in the agriculture sector, removing this source of substitutability.

The differing importance of nonhomotheticities is also intuitive. In the final consumption expenditure specification, for example, it is natural to think that food from supermarkets is a necessity, thereby leading to a negative value for $\bar{c}_a$. Similarly, it is natural to think that many services such as restaurant meals are more of a luxury, thereby leading to a positive value for $\bar{c}_s$. However, this reasoning does not apply to the consumption value added specification, since the category labeled agriculture now contains the agricultural inputs that went both into the production of “necessary” food and “luxury” restaurant meals. It follows that the nonhomotheticities should be less apparent in the consumption value added specification.

The above discussion is somewhat reminiscent of the work of Lancaster (1966). He argued that preferences should be specified over the primitive characteristics that consumers value, with the idea that individual goods are themselves bundles of characteristics. This is potentially important if we think that the set of characteristics in a given good may change over time. For example, processed food is a complex bundle of characteristics, and many of them have

changed over time. If preferences are really over characteristics and the set of characteristics in various goods are changing over time, this may manifest itself as changes in preferences if we specify them over goods instead of characteristics. While both of our specifications are at a very high level of aggregation relative to what Lancaster proposed, one might suspect that the consumption value added specification would be more stable over longer periods because it has less bundling than the final consumption expenditure specification.

Be that as it may, a key message of our calibration results is that the modeler must be very explicit about how the arguments of the utility function are being defined, since different choices imply very different parameter values. It follows that researchers must be very careful when comparing results across studies that implicitly use different definitions of the arguments in the utility function in the context of studying structural transformation. This in turn has important implications for the practice of importing parameter estimates across studies. For example, one cannot use data sets such as the Consumption Expenditure Survey to calibrate utility function parameters in a model that is defining commodities in terms of value added.

Although we have focused exclusively on preferences in our analysis, there are some obvious conclusions to be drawn regarding the role of technological factors in the process of structural transformation. In particular, given that the relative quantities of consumption value added stayed roughly unchanged while the expenditure shares changed considerably, structural transformation will require a force that changes relative prices of value added at the sector level. Examples are technological progress that is uneven at the sector level as suggested by Ngai and Pissarides (2007) and differences in capital shares at the sector level, as suggested by Acemoglu and Guerrieri (2008). In future work, we plan to quantify the relative importance of each of these forces for changes in relative prices.

5.2 Relative Merits of the Two Specifications

Each of the two specifications that we have considered has some relative merits over the alternative. One key issue is data availability. The relative advantage of the final consumption expenditure specification is that data on final consumption expenditure by category are readily available, not only from individual country sources but also in commonly used cross-country

data sets such as the Penn World Table, which measure final consumption expenditure, as opposed to production. In contrast, consumption value added data is not readily available. To be sure, data on production value added by sector is readily available, but as we argued above, this is not the same as consumption value added.

When one thinks about integrating the consumer analysis into a general equilibrium setting that has a production side, then the relative merits are reversed. If, on the one hand, the arguments of the utility function are consumption value added, then one can include production in a consistent fashion by assuming value added production functions at the sector level. These, in turn, are easily connected to data since data on value added by sector are readily available. If, on the other hand, the arguments of the utility function are final consumption expenditure across categories of goods and services, then one either needs to write down a production structure that captures the complexities of the input–output relationships at the sector level, or one needs to find a representation of production that isolates the contribution of capital and labor to the production of final expenditure categories. While this can be done, it is more difficult than working directly with sectoral value added production functions.¹⁸

Additionally, to the extent that one desires utility functions that are consistent with aggregation, given our estimates, the consumption value added specification is preferable. This is due to the relative unimportance of nonhomotheticities in this case, implying that the utility specification for consumption value added aggregates for a larger set of individual household consumption expenditure than the utility specification for final consumption expenditure.¹⁹ Moreover, the homothetic specification from Column (2) of Table 3 still provided a very good fit in the case of consumption value added, and in this case aggregation always holds.

5.3 Additional Measurement Issues

In this subsection we note five measurement issues that we have abstracted from, but which we believe should be mentioned as qualifications to our results.

The first issue is that government services are often badly measured (e.g., because value

¹⁸Valentinyi and Herrendorf (2008) showed how to construct sectoral production functions that use capital and labor to produce final expenditure by broad category.

¹⁹Formally, condition (1) in Appendix A holds for larger set of C_n .

Table 5: Decomposition of Increase in Expenditure Share of Services Consumption Value Added (accumulated 1947–2007)

Category	%
Finance, Insurance, Real Estate, Rental, and Leasing	48.8
Professional and Business Services	41.5
Health Care and Social Assistance	26.3
Information	8.3
Utilities	3.4
Educational Services	2.9
Government	1.0
Arts, Entertainment, Recreation, Accommodation, Food Services, and Other	−0.5
Trade and Transport	−31.7
	100

added is “approximated” by the corresponding wage bill). One may therefore wonder to what extent our estimation results are driven by the behavior of badly measured government services. To address this point, we have redone all of the estimations without government services.²⁰ Naturally, this reduces the quantity of services consumed, and so it lowers the estimates of the relative weight on services and the nonhomotheticity term $\bar{c}_s$. But the important question is what happens to the estimates of the elasticity of substitution and the nonhomotheticity term for agriculture. Using consumption value added data, we find that these results are literally unchanged, while using final consumption expenditure data the results are affected only somewhat.²¹ We conclude from this exercise that our results are not chiefly driven by the behavior of badly measured government services.

Second, and related, an important issue when examining time series changes in prices and quantities is the extent to which the data take proper account of quality improvements. Failure to do so will bias the decomposition of expenditure into price and quantity components. A key limitation of the official data that we have used in our analysis is that effectively no corrections are made to allow for quality improvements in services. While this is a common problem that we cannot do anything about, we think that it is worth keeping in mind.

²⁰Our initial results implicitly assumed that households were purchasing government services at the price p_s . In contrast, here we remove government services and implicitly assume that whatever utility individuals obtain from these services does not affect the marginal rate of substitution between categories of private consumption.

²¹The precise results can be found in Appendix D.

Third, consumption of durable goods typically does not equal expenditure on durable goods. For housing, which is by far the most prominent example of durables, the BEA takes account of this and imputes the rents for owner-occupied houses. For all other durables, the BEA reports expenditure (or value added) only, which forces us to associate the expenditures on these durables with current consumption. This implies, for example, that current period utility from automobiles is derived solely from current period sales (or production) of automobiles, and so we do not attribute any current period utility flow to the stock of automobiles purchased in previous periods. Because we are focused on longer term trends in aggregate data, this is not likely to be as serious as it would be in looking at individual data, or business cycle fluctuations; but it is an issue worth noting.

Forth, our model has abstracted from home production. Aguiar and Hurst (2007) and Ramey and Francis (2009) both documented a sharp drop in time devoted to home production associated with the dramatic increase in the participation rate of married women. To the extent that this has led to a substitution away from home produced services toward market produced services, our data may reflect an upward bias in the extent of the increase in the expenditure share of services.

Fifth, and related, one issue with sectoral data is the possibility that reallocation of resources across sectors reflects a relabeling of activity due to outsourcing, as opposed to fundamental shifts of economic activity across sectors. For example, if a car manufacturer changes from having in-house security guards at its establishments to purchasing security services from an outside firm, the data will record this as a movement of value added across sectors.²² This phenomenon will bias the measurement of changes in the expenditure shares of consumption value added. However, this bias is not likely to be a major driving force of structural transformation. The main reason is that industry classifications are done at the establishment level, implying that all in-house services provided at a central administrative office (headquarters) or a separate service-providing unit are always classified as service industries.

There are two additional ways of establishing that outsourcing is not the major force behind structural transformation. First, Table 5 decomposes the accumulated increase in the expen-

²²Fuchs (1968) suggested that this is one of the driving forces behind the process of structural transformation.

diture share of service consumption value added into the contributions of ten subcategories of services, where outsourced services are part of the subcategory Professional and Business Services. Although this category is the second biggest contributor to the overall increase in the expenditure share of services, more than half of the increase is accounted for by other categories. Moreover, it is reasonable to think that a substantial share of the increase in business and professional services reflects purchases directly made by consumers, in which case they would not be subject to outsourcing. A second way of establishing that outsourcing is not the major force behind structural transformation is to look at what happened to final consumption expenditure, instead of consumption value added, as final consumption expenditure are not affected by outsourcing. To stay with the example of the car manufacturer, all that matters with final consumption expenditure is how much is spent on purchases of cars. Holding the price and quantity of security services fixed, it does not matter if the security services that are implicitly reflected in the price of cars were supplied in-house or outsourced. The fact that the changes in the shares are very evident in the final consumption expenditure data confirms that the process of structural transformation is not mainly a process of outsourcing.

6 Conclusion

What utility function should one use in applied work on structural transformation and related issues? This paper sought to provide an answer to this simple question by examining the behavior of household expenditure shares for the US economy over the period 1947 to 2007. The first contribution of this paper is to clarify that given common practice in specifying multi-sector general equilibrium models, this question requires two answers, one each for two different methods of defining commodities in such models.

The second contribution of this paper is to supply the two answers. A priori there is little guidance as to how different (or similar) the two answers might be. It is very noteworthy that we find the answers to be dramatically different in terms of their basic properties. Interestingly, each of the answers takes a very simple form. If one adopts the final consumption expenditure specification, then a Stone-Geary utility function provides a very good fit to the US time se-

ries data. If instead one adopts the consumption value added specification, then a homothetic specification of the Leontief type provides a very good fit to the data.

A third contribution of the paper, which is a necessary intermediate input into the calibration of the utility function based on consumption value added, was to develop and execute a procedure for producing time series data on consumption value added. This requires extracting the component of total value added by sector that corresponds to consumption value added.

While the utility functions that we calibrate are specifically relevant for models of structural transformation, some of the basic messages of the analysis are much more general. Specifically, researchers must be careful to apply consistent definitions of commodities on both the household and production sides when connecting models with data in any multi-sector general equilibrium analysis. Changing the true definition of what is meant, for example, by the label “services” has implications not only on the household side for what form of utility function is appropriate, but also on the production side for such things as the measurement of productivity growth. This has important implications for comparing results across studies and for the practice of importing parameter values across studies.

There are several dimensions along which it will be important to extend the analysis carried out here. First, in this paper we have only analyzed the evolution of expenditure shares and prices in one country – the postwar U.S. It is also of interest to extend this analysis to a larger set of countries, in particular to situations which feature a larger range of real incomes. This will be useful in assessing the extent to which one can account for the process of structural transformation with stable preferences. Second, the work reported here focused solely on the household side. The next step in uncovering the key forces that underlie structural transformation will be to integrate this analysis with an analysis of the production side of the economy.

References

- Acemoglu, Daron and Veronica Guerrieri**, “Capital Deepening and Non–Balanced Economic Growth,” *Journal of Political Economy*, 2008, 116, 467–498.
- Aguiar, Mark and Eric Hurst**, “Measuring Leisure: The Allocation of Time Over Five Decades,” *Quarterly Journal of Economics*, 2007, 122, 969–1006.
- Atkeson, Andrew and Masao Ogaki**, “Wealth–varying Intertemporal Elasticity of Substitution: Evidence from Panel and Aggregate Data,” *Journal of Monetary Economics*, 1996, 38, 507–534.
- Bah, El-Hadj**, “Structural Transformation in Developed and Developing Countries,” manuscript, Arizona State University., Tempe, AZ 2008.
- Barnett, William A. and Apostolos Serletis**, “Consumer Preferences and Demand Systems,” *Journal of Econometrics*, 2008, 147, 210–224.
- Baumol, William J.**, “Macroeconomics of Unbalanced Growth: The Anatomy of the Urban Crisis,” *American Economic Review*, 1967, 57, 415–426.
- Blundell, Richard**, “Consumer Behaviour: Theory and Evidence – A Survey,” *Economic Journal*, 1988, 98, 16–65.
- Buera, Francisco J. and Joseph P. Kaboski**, “Can Traditional Theories of Structural Change Fit The Data?,” *Journal of the European Economic Association*, 2009, 7, 469–477.
- Bureau of Economic Analysis**, “Concepts and Methods of the U.S. Input–Output Accounts,” 2006. http://bea.gov/bea/papers/IOmanual_092906.pdf.
- Caselli, Francesco and Wilbur John Coleman**, “The U.S. Structural Transformation and Regional Convergence: A Reinterpretation,” *Journal of Political Economy*, 2001, 109, 584–616.
- Deaton, Angus and John Muellbauer**, “An Almost Ideal Demand System,” *American Economic Review*, 1980, 70, 312–326.

- Duarte, Margarida and Diego Restuccia**, “The Role of the Structural Transformation in Aggregate Productivity,” *Forthcoming: Quarterly Journal of Economics*, 2009.
- Echevarria, Cristina**, “Changes in Sectoral Composition Associated with Economic Growth,” *International Economic Review*, 1997, 38, 431–452.
- Foellmi, Reto and Josef Zweimuller**, “Structural Transformation in Developed and Developing Countries,” *Journal of Monetary Economics*, 2008, 55, 1317–1328.
- Fuchs, Victor**, *The Service Economy*, New York: Columbia University Press, 1968.
- Geary, Roy C.**, “A Note on A Constant–Utility Index of the Cost of Living,” *Review of Economic Studies*, 1950–1951, 18, 65–66.
- Gollin, Douglas, Stephen L. Parente, and Richard Rogerson**, “The Role of Agriculture in Development,” *American Economic Review, Papers and Proceedings*, 2002, 92, 160–164.
- Herrendorf, Berthold, James Schmitz Jr, and Arilton Teixeira**, “Transportation and Development: Insights from the US, 1840–1860,” Staff Report 425, Federal Reserve Bank of Minneapolis 2009.
- Houthakker, Hendrik S. and Lester Taylor**, *Consumer Demand in the United States: Analysis and Projections*, Cambridge, MA: Harvard University Press, 1970.
- Huffman, Gregory W. and Mark A. Wynne**, “The Role of Intratemporal Adjustment Costs in a Multisector Economy,” *Journal of Monetary Economics*, 1999, 43, 317–350.
- Kongsamut, Piyabha, Sergio Rebelo, and Danyang Xie**, “Beyond Balanced Growth,” *Review of Economic Studies*, 2001, 68, 869–882.
- Kuznets, Simon**, *Modern Economic Growth*, New Haven: Yale University Press, 1966.
- Laitner, John**, “Structural Change and Economic Growth,” *Review of Economic Studies*, 2000, 67, 545–561.
- Lancaster, Kelvin**, “A New Approach to Consumer Theory,” *Journal of Political Economy*, 1966, 74, 132–157.

- Landefeld, J. Steven and Robert P. Parker**, “BEA’s Chain Indexes, Time Series, and Measures of Long Term Economic Growth,” in “Survey of Current Business,” Bureau of Economic Analysis, May 1997.
- Messina, Julian**, “The Role of Product Market Regulations in the Process of Structural Change,” *European Economic Review*, 2006, 50, 1863–1890.
- Ngai, L. Rachel and Christopher A. Pissarides**, “Structural Change in a Multisector Model of Growth,” *American Economic Review*, 2007, 97, 429–443.
- **and —**, “Trends in Hours and Economic Growth,” *Review of Economic Dynamics*, 2008, 11, 239–256.
- Pollak, Robert and Terence J. Wales**, “Estimation of the Linear Expenditure System,” *Econometrica*, 1969, 37, 611–628.
- Ramey, Valery A. and Neville Francis**, “A Century of Work and Leisure,” *AEJ: Macro*, 2009, 1, 189–224.
- Rogerson, Richard**, “Structural Transformation and the Deterioration of European Labor Market Outcomes,” *Journal of Political Economy*, 2008, 226, 235–259.
- Sato, Ryuzo**, “The Most General Class of CES Functions,” *Econometrica*, 1975, 43, 999–1003.
- Stone, Richard**, “Linear Expenditure Systems and Demand Analysis: An Application to the Pattern of British Demand,” *Economic Journal*, 1954, 64, 511–527.
- ten Raa, Thijs**, *The Economics of Input-Output Analysis*, Cambridge, UK: Cambridge University Press, 2005.
- Valentinyi, Akos and Berthold Herrendorf**, “Measuring Factor Income Shares at the Sectoral Level,” *Review of Economic Dynamics*, 2008, 11, 820–835.

Appendix A: Aggregation of Demand Functions

Consider N households indexed by $n = 1, \dots, N$. Each household solves:

$$\begin{aligned} \max_{c_a^n, c_m^n, c_s^n} & \left[\omega_a^\sigma (c_a^n + \bar{c}_a)^{\frac{\sigma-1}{\sigma}} + \omega_m^\sigma (c_m^n + \bar{c}_m)^{\frac{\sigma-1}{\sigma}} + \omega_s^\sigma (c_s^n + \bar{c}_s)^{\frac{\sigma-1}{\sigma}} \right]^{\frac{\sigma}{\sigma-1}} \\ \text{s.t.} \quad & p_a c_a^n + p_m c_m^n + p_s c_s^n \leq C_n. \end{aligned}$$

Let the parameters and the income distribution be such that for all $n \in \{1, \dots, N\}$ household expenditure exceed a minimum level:

$$C_n > \sum_{i=a,m,s} p_i \max\{-\bar{c}_i, 0\}. \quad (1)$$

Then the solution to each household's problem is interior and the first-order conditions are

$$\begin{aligned} \left(\frac{\omega_a}{\omega_s} \right)^{\frac{1}{\sigma}} \left(\frac{c_a^n + \bar{c}_a}{c_s^n + \bar{c}_s} \right)^{-\frac{1}{\sigma}} &= \frac{p_a}{p_s}, \\ \left(\frac{\omega_m}{\omega_s} \right)^{\frac{1}{\sigma}} \left(\frac{c_m^n + \bar{c}_m}{c_s^n + \bar{c}_s} \right)^{-\frac{1}{\sigma}} &= \frac{p_m}{p_s}, \end{aligned}$$

which can be rewritten as

$$\begin{aligned} \frac{p_a}{p_s} \frac{c_a^n + \bar{c}_a}{c_s^n + \bar{c}_s} &= \frac{\omega_a}{\omega_s} \left(\frac{p_a}{p_s} \right)^{1-\sigma}, \\ \frac{p_m}{p_s} \frac{c_m^n + \bar{c}_m}{c_s^n + \bar{c}_s} &= \frac{\omega_m}{\omega_s} \left(\frac{p_m}{p_s} \right)^{1-\sigma}. \end{aligned}$$

This gives the demand functions

$$\begin{aligned}
p_a(c_a^n + \bar{c}_a) &= \frac{p_a(c_a^n + \bar{c}_a) + p_m(c_m^n + \bar{c}_m) + p_s(c_s^n + \bar{c}_s)}{1 + \frac{\omega_m}{\omega_a} \left(\frac{p_m}{p_a} \right)^{1-\sigma} + \frac{\omega_s}{\omega_a} \left(\frac{p_s}{p_a} \right)^{1-\sigma}}, \\
p_m(c_m^n + \bar{c}_m) &= \frac{p_a(c_a^n + \bar{c}_a) + p_m(c_m^n + \bar{c}_m) + p_s(c_s^n + \bar{c}_s)}{1 + \frac{\omega_a}{\omega_m} \left(\frac{p_a}{p_m} \right)^{1-\sigma} + \frac{\omega_s}{\omega_m} \left(\frac{p_s}{p_m} \right)^{1-\sigma}}, \\
p_s(c_s^n + \bar{c}_s) &= \frac{p_a(c_a^n + \bar{c}_a) + p_m(c_m^n + \bar{c}_m) + p_s(c_s^n + \bar{c}_s)}{1 + \frac{\omega_a}{\omega_s} \left(\frac{p_a}{p_s} \right)^{1-\sigma} + \frac{\omega_m}{\omega_s} \left(\frac{p_m}{p_s} \right)^{1-\sigma}}.
\end{aligned}$$

Adding up over all households, we obtain:

$$\begin{aligned}
p_a(c_a + N\bar{c}_a) &= \frac{p_a(c_a + N\bar{c}_a) + p_m(c_m + N\bar{c}_m) + p_s(c_s + N\bar{c}_s)}{1 + \frac{\omega_m}{\omega_a} \left(\frac{p_m}{p_a} \right)^{1-\sigma} + \frac{\omega_s}{\omega_a} \left(\frac{p_s}{p_a} \right)^{1-\sigma}}, \\
p_m(c_m + N\bar{c}_m) &= \frac{p_a(c_a + N\bar{c}_a) + p_m(c_m + N\bar{c}_m) + p_s(c_s + N\bar{c}_s)}{1 + \frac{\omega_a}{\omega_m} \left(\frac{p_a}{p_m} \right)^{1-\sigma} + \frac{\omega_s}{\omega_m} \left(\frac{p_s}{p_m} \right)^{1-\sigma}}, \\
p_s(c_s + N\bar{c}_s) &= \frac{p_a(c_a + N\bar{c}_a) + p_m(c_m + N\bar{c}_m) + p_s(c_s + N\bar{c}_s)}{1 + \frac{\omega_a}{\omega_s} \left(\frac{p_a}{p_s} \right)^{1-\sigma} + \frac{\omega_m}{\omega_s} \left(\frac{p_m}{p_s} \right)^{1-\sigma}},
\end{aligned}$$

where

$$c_i \equiv \sum_{n=1}^N c_i^n.$$

Let $C \equiv \sum_{n=1}^N C^n$. If the stand-in household solves

$$\begin{aligned}
&\max_{c_a, c_m, c_s} \left[\omega_a(c_a + N\bar{c}_a)^{\frac{\sigma-1}{\sigma}} + \omega_m(c_m + N\bar{c}_m)^{\frac{\sigma-1}{\sigma}} + \omega_s(c_s + N\bar{c}_s)^{\frac{\sigma-1}{\sigma}} \right]^{\frac{\sigma}{\sigma-1}} \\
&\text{s.t. } p_a c_a + p_m c_m + p_s c_s \leq C,
\end{aligned}$$

then its choices satisfy

$$c_i = \sum_{n=1}^N c_i^n.$$

In other words, there is aggregation.

Appendix B: Data Sources and Sector Assignment

B.1: Data Sources

All data are for the U.S. during 1947–2007.

We calculate a per capita quantity by dividing the total quantity by the population size. We take the population size from Table 7.1: “Selected Per Capita Product and Income Series in Current and Chained Dollars”.

The construction of final consumption expenditure data is based on standard NIPA tables from the BEA. We use the most recent NIPA data released in August 2009 which incorporates the last comprehensive revision. In particular, we use data from the following tables:

- Table 2.4.3: “Real Personal Consumption Expenditures by Type of Product, Quantity Indexes”; Table 2.4.5: “Personal Consumption Expenditures by Type of Product”.
- Table 3.10.3: “Real Government Consumption Expenditures and General Government Gross Output, Quantity Indexes”; Table 3.10.5: “Government Consumption Expenditures and General Government Gross Output”

The construction of production value added data by sector is based on the Annual Industry Accounts, which contain current dollar value added and quantity indices by industry based on chain weighted methods. The value added by industry data is consistent with the NAICS for the entire period 1947–2007.²³

The construction of consumption value added (as opposed to production value added) is based on two main data sources: the annual expenditure data described above and the total

²³<http://www.bea.gov/industry/gpotables/AllTables.zip>

requirement matrices from the IO Tables. In the next subsection, we describe in detail how these two data sources are combined to obtain consumption value added. Here we just describe the exact data sources. There are benchmark IO Tables and annual IO Tables. Benchmark IO Tables are available for 1947, 1958, 1963, 1967, 1972, 1977, 1982, 1987, 1992, 1997 and 2002.²⁴ Annual IO Tables are available for each year during the period 1998–2007.²⁵ An important additional data source are the so called “Bridge Tables for Personal Consumption Expenditure”, which are available for the 1997 and 2002 benchmark IO Tables. Bridge Tables link IO Tables with the standard expenditure data of the BEA. In particular, they report how personal consumption expenditure in the IO Tables are related to those in the BEA expenditure tables. If we don’t have IO Tables for a particular year, then we use linear interpolation between the years for which IO Tables are available.

B.2: Sector Assignment

When we use final consumption expenditure data, the three sectors contain the following BEA *commodities*:

- Agriculture: “food and beverages purchased for off-premises consumption”
- Manufacturing: “durable goods”; “nondurable goods” excluding “food and beverages purchased for off-premises consumption”
- Services: “services”; “government consumption expenditure”

When we use value added data, the three sectors contain the following BEA *industries*:

- Agriculture: “farms”; “forestry, fishing, and related activities”
- Manufacturing: “construction”; “manufacturing”; “mining”
- Services: all other industries including “government industries”

²⁴http://www.bea.gov/industry/io_benchmark.htm

²⁵http://www.bea.gov/industry/io_annual.htm

Appendix C: Calculating Consumption Value Added

C.1: Constructing Final Expenditure in Producer's Prices

C.1.1: Disaggregation to six sectors

To obtain final consumption expenditure in producer's prices from the available data on final consumption expenditure in purchaser's prices, we need to remove the distribution costs from the different goods categories and move them to services. For two reasons, this requires further disaggregation. First, we calculate the distribution costs for retail, wholesale and transportation services from the expenditure on the sector Trade and Transport. We therefore, need to separate Trade and Transport from the rest of services. Second, the expenditure on mining involve distribution costs whereas those on construction do not, so we need to separate the two from other manufacturing. We therefore consider the following six sectors: Agriculture, Mining, Construction, Manufacturing, Trade and Transport, and Services excluding Trade and Transport, which we index by $i \in \{Ag, Mi, Co, Ma, TT, Se\}$, which aggregate to our model sectors in the obvious way: $a = \{Ag\}$, $m = \{Mi, Co, Ma\}$, $s = \{TT, Se\}$. Note that while we use the BEA classification for Agriculture, Mining, Construction, and Manufacturing, the sector Trade and Transport combines "Wholesale Trade", "Retail Trade" and "Transportation and Warehousing".

We should mention a potential problem that arises from the reclassification of industries over time. In particular, while the BEA now publishes GDP by industry data based on the NAICS for the whole period 1947–2007, it still publishes the underlying input–output tables for the subperiod 1947–1977 based on the different SIC's. Fortunately, many of the reclassifications from the SIC's to the NAICS happened at finer levels of disaggregation than we study here, and so they do not affect the aggregates of the six sectors we have just introduced. However, there are some exceptions that we need to reclassify to make the input–output tables consistent with the GDP by industry data. The most important example is the "Publishing Industries", which the SIC has as manufacturing industry and the NAICS has as a service industry.

C.1.2: Removing distribution costs from personal consumption expenditure

We now explain how to remove distribution costs from personal consumption expenditure.

The expenditure side of GDP values personal consumption expenditure at purchaser's prices and it disaggregates them into the expenditure on goods, trade and transportation, and services excluding trade and transportation:

$$PC^{Pu} = PC_{Gs}^{Pu} + PC_{TT}^{Pu} + PC_{Se}^{Pu}.$$

Goods consist of “durable and nondurable goods” excluding “food and beverages purchased for off-premises consumption”, trade and transportation consists of “public transportation”, and services consist of “services” excluding “public transportation”.

We start by removing distribution costs from personal consumption expenditure on goods. To go from purchaser's to producer's prices, we calculate the distribution margins $DM_{PC_{Gs}}$ by using the fact that in the IO Tables personal consumption expenditure on trade and transportation consists of all transportation expenditure whereas PC_{TT}^{Pu} consists only of “public transportation” that households explicitly purchase. Hence, the difference between the two equals the distribution costs of goods that household purchase indirectly when purchasing goods, and so:

$$DM_{PC_{Gs}} = \frac{(PC_{TT}^{IO} - PC_{TT}^{Pu})}{(PC_{TT}^{IO} - PC_{TT}^{Pu}) + (PC_{Ag}^{IO} + PC_{Mi}^{IO} + PC_{Co}^{IO} + PC_{Ma}^{IO})},$$

$$PC_{Gs}^{Pr} = (1 - DM_{PC_{Gs}})PC_{Gs}^{Pu}.$$

We continue by removing distribution costs from personal consumption expenditure on services. This is straightforward because the IO Tables suggest that personal consumption expenditure on services involve negligible distribution costs. Therefore:

$$PC_{Se}^{Pr} = PC_{Se}^{Pu}.$$

Given that we have calculated PC_{Gs}^{Pr} , we now disaggregate it into the components PC_{Ag}^{Pr} , PC_{Mi}^{Pr} , PC_{Co}^{Pr} and PC_{Ma}^{Pr} . The IO Tables report that PC_{Mi}^{IO} are very small and that PC_{Co}^{IO} are zero in all years. We therefore set $PC_{Mi}^{Pr} = PC_{Co}^{Pr} = 0$. This leaves us with the task of splitting PC_{Gs}^{Pr} between PC_{Ag}^{Pr} and PC_{Ma}^{Pr} . First, we calculate expenditures on food at producer prices, PC_{Food}^{Pr} .

Expenditure on food is “food and beverages purchased for off-premises consumption”. We remove distribution costs by applying the distribution margin of goods that we calculated above, $PC_{Food}^{Pr} = (1 - DM_{Gs}^{PC})PC_{Food}^{Pu}$. Next, since PC_{Food}^{Pr} contains both unprocessed and processed food, we need to take processed food out to obtain the expenditure on agricultural commodities. We use that PC_{Ag}^{IO} are the expenditure on agricultural goods without processed food. Defining $\Phi_1 \equiv PC_{Ag}^{IO}/PC_{Food}^{Pr}$, we have

$$PC_{Ag}^{Pr} = \Phi_1 PC_{Food}^{Pr}.$$

In sum, the components of personal consumption expenditure in producer’s prices are obtained as follows:

$$\begin{aligned} PC_{Ag}^{Pr} &= \Phi_1 PC_{Food}^{Pr} \\ PC_{Mi}^{Pr} &= 0, \\ PC_{Co}^{Pr} &= 0, \\ PC_{Ma}^{Pr} &= (1 - DM_{PCGs})PC_{Gs}^{Pu} - PC_{Ag}^{Pr}, \\ PC_{TT}^{Pr} &= PC_{TT}^{Pu} + DM_{PCGs}PC_{Gs}^{Pu}, \\ PC_{Se}^{Pr} &= PC_{Se}^{Pu}. \end{aligned}$$

C.1.3: Removing distribution costs from government consumption expenditure

We now explain how to remove distribution costs from final expenditure on government consumption.

In the IO Tables, the general government appears as a production industry and as a commodity. In the expenditure side of GDP, government consumption expenditure at purchaser’s prices are defined as the gross output of the general government industry minus own account investment and sales to other sectors.

The treatment of the gross output of the general government industry changed in 1998. Before 1998, it was defined as its value added GC_{VA}^{Pu} (compensation of general government employees plus consumption of general government fixed capital). All intermediate inputs

were consequently treated as final government expenditure on these goods. Since 1998, the gross output of the general government industry has included intermediate goods, that is, it is defined as the sum of value added GC_{VA}^{Pu} , purchased intermediate goods, GC_{Gs}^{Pu} , and purchased intermediate services, GC_{Se}^{Pu} .

We start with the period 1947–1997. During this period, the IO Tables show that GC_{Ag}^{IO} and GC_{Mi}^{IO} are small, so we set $GC_{Ag}^{Pr} = GC_{Mi}^{Pr} = 0$. The distribution margins of government consumption expenditure in the 1997 IO Tables on average equal 18% of the distribution margins of personal consumption expenditure, so we set $DM_{GC_{Gs}} = 0.18 \cdot DM_{PC_{Gs}}$. Next we calculate GC_{Co}^{Pr} . The raw IO Tables distinguish between government expenditure on “maintenance and repair construction” and on “new construction”. First, we calculate

$$\Phi_2 \equiv \frac{\text{government expenditure on maintenance and repair construction}}{\text{depreciation on government structures}},$$

where the depreciation on government structures is taken from Table 7.3: “Current-Cost Depreciation of Government Fixed Assets” of the BEA Fixed Assets Tables. We then calculate GC_{Co}^{Pr} by multiplying Φ_2 with depreciation on government structures.

In sum, for the period 1947–1997, we calculate the variables of interest as:

$$GC_{Ag}^{Pr} = 0, \tag{2a}$$

$$GC_{Mi}^{Pr} = 0, \tag{2b}$$

$$GC_{Co}^{Pr} = \Phi_2 \text{Depreciation on government structures}, \tag{2c}$$

$$GC_{Ma}^{Pr} = (1 - DM_{GC_{Gs}})GC_{Gs}^{Pu} - GC_{Co}^{Pr}, \tag{2d}$$

$$GC_{TT}^{Pr} = DM_{GC_{Gs}}GC_{Gs}^{Pu}, \tag{2e}$$

$$GC_{Se}^{Pr} = GC_{VA}^{Pu} + GC_{Se}^{Pu} - \text{Sales to other sectors} - \text{Own account investment}. \tag{2f}$$

The rationale behind the fourth equation is that own account investment typically involves goods, so it has to be taken out of government intermediate good expenditure. The fifth equation is just an implication of the fourth equation. The last equation expresses that government expenditure on services are equal to the value added representing the service flow from gov-

ernment capital and employees plus the services purchased as intermediate input net of what is invested on own account sold to other sectors, which typically are general government services. The equation reflects that typically services do not have distribution costs, so services evaluated at producer's and purchaser's prices are the same.

For the period 1998–2007, government consumption expenditure in the IO Tables almost exclusively consist of expenditure on general government services. Since services have no distribution costs, we set:

$$GC_{Ag}^{Pr} = GC_{Mi}^{Pr} = GC_{Co}^{Pr} = GC_{Ma}^{Pr} = GC_{TT}^{Pr} = 0,$$

$$GC^{Pr} = GC^{Pu}.$$

C.2: Linking Consumption Expenditures to Value Added

The total requirement matrix (henceforth TR Matrix) links the income and the expenditure side of GDP. We now explain how to use the TR Matrix to obtain the value added in producer's prices that are generated by the final expenditure on consumption in producer's prices, which we have just constructed in the previous subsection. We use the language and the notation of the BEA to the extent possible. For further explanation see ten Raa (2005) and Bureau of Economic Analysis (2006).

The way in which the TR Matrix is calculated changed in 1972. So for years prior to 1972, the IO Tables assumed that each industry produces one commodity and that each commodity is produced in exactly one industry. For years after 1972 the IO Tables have taken account of the fact that industries can produce more than one commodity and that the same commodity can be produced in different industries.

We start by explaining the TR Matrix prior to 1972. We denote the number of industries by n , which before 1972 equals the number of commodities. Domestically produced commodities are purchased either by domestic industries (intermediate expenditure) or by final users (final uses or final expenditure). Final uses include both domestic final uses and exports, where exports can be either intermediate or final foreign uses. Domestic industries produce gross output and the difference between gross output and intermediate expenditure is industry value

added.

Let $\mathbf{A}$ denote the $(n \times n)$ transaction matrix.²⁶ Rows are associated with commodities and columns with industries: entry ij shows the dollar amount of commodity i that industries j uses per dollar of output it produces. Note that these commodities may have been produced domestically or imported. Let $\mathbf{q}$ denote the $(n \times 1)$ output vector of domestically produced commodities. Element i records the sum of the dollar amounts of commodity i that are delivered to other domestic industries as intermediate inputs and to final uses. Let $\mathbf{g}$ denote the $(n \times 1)$ industry output vector. Element j records the dollar amount of output of industry j . Let $\mathbf{e}$ denote the $(n \times 1)$ vector of expenditures on final uses. Element i records the dollar amount of final uses of the domestically produced commodity i , so component e_i reports domestic private and public consumption, domestic investment, and net exports of commodity i .

Two identities link these vectors and with the TR Matrix:

$$\mathbf{q} = \mathbf{A}\mathbf{g} + \mathbf{e}, \quad (3)$$

$$\mathbf{q} = \mathbf{g}. \quad (4)$$

The first identity states that the dollar amount of domestically produced output of each commodity equals the sum of intermediate uses plus the final uses of that commodity. The second identity states that total value of output of industry i equals to the total value of commodity i , which is trivially true here because each industry is assumed to produce one distinct commodity. We can solve these two equations for $\mathbf{g}$:

$$\mathbf{g} = (\mathbf{I} - \mathbf{A})^{-1}\mathbf{e}, \quad (5)$$

where $\mathbf{I}$ is the $(n \times n)$ identity matrix (1 in the diagonal and zero elsewhere). $\mathbf{R} \equiv (\mathbf{I} - \mathbf{A})^{-1}$ is called the total requirements matrix. Rows are associated with industries and columns with commodities. Entry ji shows the dollar value of industry j 's production that is required, both directly and indirectly, to deliver one dollar of the domestically produced commodity i to final uses including net exports.

²⁶Matrices and vectors are in bold symbol throughout the paper.

We continue by explaining the TR Matrix after 1972, so now the IO Tables take account of the fact that an industry may produce more than one commodity and that a commodity may be produced in different industries. In general, the number of industries n will then differ from the number of commodities. We call the number of commodities m . This implies that we no longer have one transaction matrix, but a use and a make matrix. $\mathbf{B}$ denotes the $(m \times n)$ use matrix. Entry ij shows the dollar amount of commodity i that industries j uses per dollar of output it produces. Again note that these commodities may have been produced domestically or imported. $\mathbf{W}$ denotes the $(n \times m)$ make matrix. Rows are associated with industries and columns with commodities: entry ji shows which share of one dollar of the domestically produced commodity i industry j makes. Two identities link these matrices and vectors:

$$\mathbf{q} = \mathbf{B}\mathbf{g} + \mathbf{e}, \quad (6a)$$

$$\mathbf{g} = \mathbf{W}\mathbf{q}. \quad (6b)$$

The first identity says that the dollar amount of each domestically produced commodity equals the sum of the dollar amount of that commodity that the different domestic industries use as intermediate goods plus the dollar amount of final uses of that commodity. Note again that final uses are for domestic private and public consumption, domestic investment, and net exports. The second identity says the dollar output of each industry equals the sum of that industry's contribution to the outputs of the different domestically produced commodities.

To eliminate $\mathbf{q}$ from these identities, we substitute (6b) into (6a) to obtain $\mathbf{q} = \mathbf{B}\mathbf{W}\mathbf{q} + \mathbf{e}$. We then solve this for $\mathbf{q}$ and substitute the result back into (6b). This gives:

$$\mathbf{g} = \mathbf{W}(\mathbf{I} - \mathbf{B}\mathbf{W})^{-1}\mathbf{e}. \quad (7)$$

$\mathbf{R} \equiv \mathbf{W}(\mathbf{I} - \mathbf{B}\mathbf{W})^{-1}$ is called the industry-by-commodity total requirements matrix. Rows are associated with industries and columns with commodities. Entry ji shows the dollar value of industry j 's production that is required, both directly and indirectly, to deliver one dollar of the domestically produced commodity i to final uses including net exports.

Let $\mathbf{v}$ denote the $(1 \times n)$ vector of industry value added per unit of industry output, which

is easily calculated from the IO Tables by dividing industry value added by industry output. To obtain the value added, $\mathbf{va}$, that is generated by the domestically produced final expenditure vector, $\mathbf{e}$, we multiply $\mathbf{R}$ (as defined either in (5) or in (7)) with $\mathbf{e}$:

$$\mathbf{va} = \langle \mathbf{v} \rangle \mathbf{R} \mathbf{e},$$

where $\langle \mathbf{v} \rangle$ denotes the diagonal matrix with vector $\mathbf{v}$ in its diagonal. It is important to realize that this formula works for any domestically produced final expenditure vector, and so in principle we could use it for the domestically produced final consumption vector. However, we don't know this vector because we do not know the share of imports that is consumed. Instead, we only know the final consumption vector, $\mathbf{e}_c$. Component i of this vector reports the final consumption of commodity i , which may either be produced domestically or be imported. Assuming that imported commodities are produced with the same input requirements as in the U.S., we can use the total requirement matrix together with the vector $\mathbf{e}_c$. This gives us the consumption value added vector we are looking for:

$$\mathbf{c} = \langle \mathbf{v} \rangle \mathbf{R} \mathbf{e}_c. \quad (8)$$

In words, the vector on the left-hand side reports the value added in the different industries that is generated by the final consumption expenditure vector $\mathbf{e}_c$. Aggregating the components of this vector into our three broad sectors agriculture, manufacturing, and services gives us the consumption value added used in the text.

Appendix D: Results Without Government Services

Here we offer the results of our calibration exercises when we exclude government services. Table 6 reports the results. For convenience, the first and the third column repeat the results with government services from Table 1 and Table 3.

Table 6: Results With and Without Government Services

	Final Consumption Expenditure		Consumption Value Added	
	With	Without	With	Without
σ	.81	.52	0	0
$\bar{c}_a$	-1,208	-1,246	-136.2	-130.0
$\bar{c}_s$	8,024	2,881	3,652	931.3
ω_m	.18	.25	.15	.17
ω_s	.81	.71	.84	.82
Loss	.105	.122	.124	.113