Volume Title: A New Architecture for the U.S. National Accounts

Volume Author/Editor: Dale W. Jorgenson, J. Steven Landefeld, and William D. Nordhaus, editors

Volume Publisher: University of Chicago Press

Volume ISBN: 0-226-41084-6

Volume URL: http://www.nber.org/books/jorg06-1

Conference Date: April 16-17, 2004

Publication Date: May 2006

Title: Integrating Expenditure and Income Data: What to Do with the Statistical Discrepancy?

Author: J. Joseph Beaulieu, Eric J. Bartelsman

URL: http://www.nber.org/chapters/c0141

Integrating Expenditure and Income Data
What to Do with the Statistical Discrepancy?

J. Joseph Beaulieu and Eric J. Bartelsman

> A man with one watch knows what time it is:
> A man with two watches is never quite sure.
> —French Proverb

8.1 Introduction

The Bureau of Economic Analysis (BEA) publishes two measures of domestic output. The better-known measure, gross domestic product (GDP), is the sum of private and government consumption and investment (including inventory investment) and net exports. A second measure, gross domestic income (GDI), is the sum of factor and nonfactor payments paid to input providers; these payments include compensation, profits and profit-like income, production and import taxes (formerly known as indirect business taxes), and the consumption of fixed capital. GDP and GDI conceptually measure the same thing, but because the two are calculated using imperfect source data, the two measures differ by what is called the statistical discrepancy.

Historically, the level of the statistical discrepancy has been small relative to GDP or GDI. As shown in panel A of figure 8.1, the absolute value of the statistical discrepancy as a fraction of the average of nominal GDP and nominal GDI peaked at 2.1 percent in 1993. From 1977 to 2001, the fraction averaged 0.8 percent with a standard deviation of 0.9 percent.

Nonetheless, different movements in real GDP and in real GDI can be economically meaningful. Panel B of figure 8.1 plots the average annual

A

The Statistical Discrepancy
(Percent of the Average of GDP and GDI, Annual Data)

B

Real GDP and Real GDI
(Percent Change of Annual Averages)
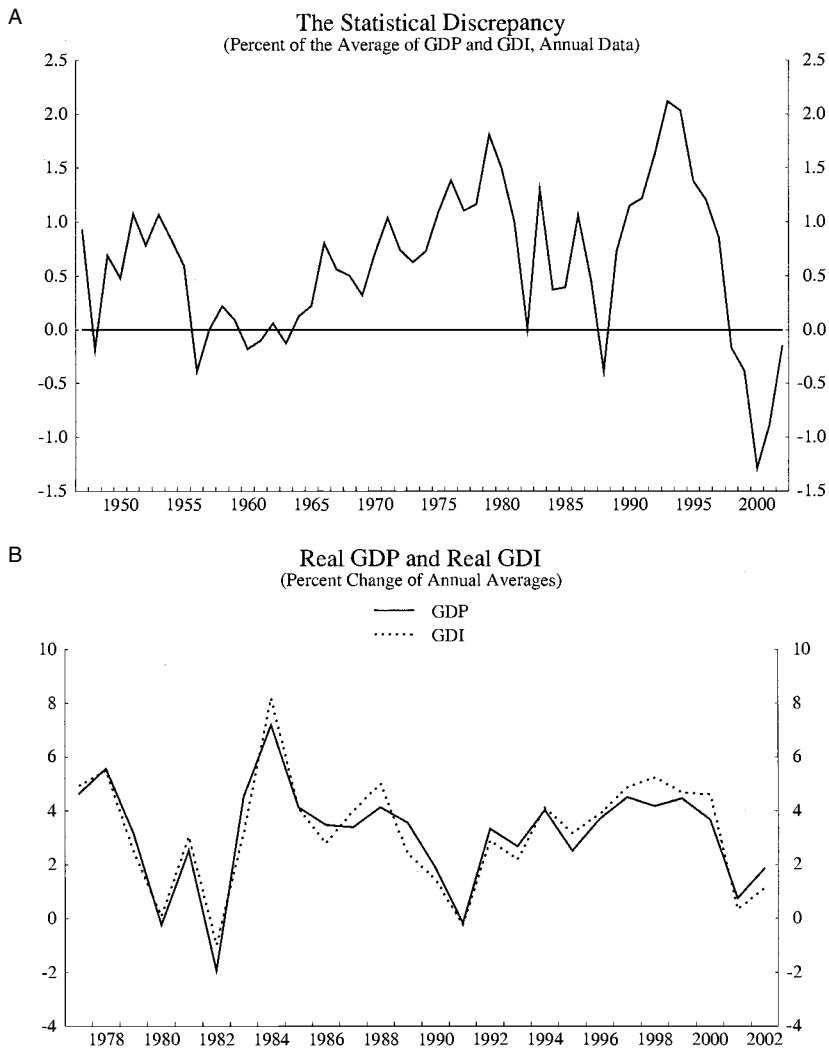
——— GDP
······ GDI

Fig. 8.1

growth rates of real GDP and GDI. Although the movements of the two appear to coincide from year to year, between 1994 and 2000, real GDI grew on average 1/2 percentage point (annual rate) faster than real GDP, which is sizable when compared to the average growth rate of the two series of 4.1 percent.

The recent difference in the growth rates of the two measures of domestic product has been a problem for policymakers. The two measures imply

different paths for productivity and potential output, which are important for planning purposes. Many analysts have pointed to the rapid rate of growth of GDI as being more consistent with the expected productivity gains from investment in high-tech equipment. Problems for analysts are especially acute when they need to combine data from the expenditure and income accounts, such as when modeling the components of national saving or projecting tax receipts. Indeed, the Congressional Budget Office (CBO) points to the large swing in the statistical discrepancy as a substantial hindrance in its ability to forecast tax revenue in the past few years (CBO 2003). The statistical discrepancy also leads to inconsistencies when analyzing particular types of income as a share of GDP.

Finally, the existence of the statistical discrepancy is a problem for researchers trying to reconcile their estimates of productivity trends by industry using data measured on the income side with aggregate estimates of productivity trends that are based on product-side measures. Bartelsman and Beaulieu (2004), Bosworth and Triplett (2003), and Nordhaus (2000) use the BEA's GDP-by-industry data (2003 or earlier) to model industry-level productivity. These data aggregate to GDI, making it hard to compare their results to the BLS's measure of productivity in the nonfarm business sector, which equals GDP less the value added from a few select sectors.[1]

Several researchers have speculated on the data deficiencies that have led to the statistical discrepancy. GDP may be mismeasured because estimating the consumption of services is difficult (Council of Economic Advisers 1997; Moulton 2000) or exports are underreported (Moulton 2000). GDI may be mismeasured because purging income of capital gains, which do not represent current production, is hard (Baker 1998; Moulton 2000), because stock options and other nontraditional forms of compensation show up in the compensation statistics without an offset in the profits data (Baker 1998; Moulton 2000), or because measures of proprietors' income have to be adjusted for underreporting in the tax return data. These adjustments to proprietors' income are based on an outdated and discontinued study (Council of Economic Advisers 1997). Many of these explanations appear to be confirmed by Klein and Makino (2000), who find that the statistical discrepancy is inversely related to profits and proprietors' income and positively related to government spending and exports.[2]

---

1. Despite what one may infer from the name "gross domestic product by industry," the industry estimates in this data set aggregate to GDI. A balancing item is included in this data set, but this discrepancy is not allocated across industries; see Yuskavage and Strassner (2003). The BEA has recently altered its methodology to produce industry data, and its latest estimates of these data now aggregate to GDP; see Lawson et al. (chap. 6 in this volume).

2. Recall the convention that more GDP relative to GDI leads to a more positive statistical discrepancy; more GDI leads to a more negative discrepancy.

The BEA prefers GDP as its measure of domestic output. Parker and Seskin (1997) write:

> [The BEA] considers the source data underlying the estimates of GDP to be more accurate. For example, most of the annual source data used for estimating GDP are based on complete enumerations, such as the Federal Government budget data, or are regularly adjusted to complete enumerations, such as the quinquennial economic censuses and census of governments. . . . For GDI, only the annual tabulations of employment tax returns and Federal Government budget data are complete enumerations, and only farm proprietors' income and State and local government budget data are regularly adjusted to complete enumerations. For most of the remaining components of GDI, the annual source data are tabulations of samples of income tax returns.

This view is reflected in the presentation of the NIPAs. The BEA presents only GDP-related data in its summary tables, and in its decomposition of national income it portrays the statistical discrepancy as if it were all an error in the measurement of income vis-à-vis GDP. A few years ago, the BLS appeared to adopt this view when it switched its definition of nonfarm business output in its Productivity and Cost release from one based on GDI to one based on GDP, as described in Dean, Harper, and Otto (1995).

Others, however, have argued that GDI has more desirable properties, at least at certain points in time. The Council of Economic Advisers (1997) found that the behavior of Okun's law, the sharp jump in personal tax payments, and the behavior of the real product wage were more consistent with the faster-growing GDI measure of output in the mid-1990s, as measured at that time. During that same period, Greenspan (2004) observed that the rapid rise in measured labor and capital income, along with quiescent price inflation, suggested that productivity was increasing briskly. These productivity gains were apparent in the income-side measure, but not in the product-side measure of domestic output. Based on their time-series properties, Weale (1992) argued that GDI should be weighted almost twice as much as GDP in an optimal combination of the two measures into a single output series.

The paper presents two sets of exercises. One is to conduct a "forensic" examination of the statistical discrepancy by allocating the statistical discrepancy across industries; perhaps we can lessen the size of the aggregate discrepancy through focused, improved measurement at the industry level. Next, we present some metrics that allow us to evaluate a sequence of data sets created under varying assumptions regarding the quality of the underlying data sources. Optimizing on these metrics should provide one, best, coherent data set to conduct further research.

The structure of the paper is as follows. In section 8.2 we describe the underlying source data, the manipulations to the data undertaken to make the sources consistent in classifications and definitions, and the method

used to integrate the varying source data. In section 8.3, we compare estimates of value added by industry from a consistent data set controlled to GDP data with value added by industry from a consistent data set controlled to GDI data to calculate statistical discrepancies by industry. Two sets of estimates of deliveries to final demand by industry also yield statistical discrepancies by industry. Similarly, we compare our two sets of estimates of final demand by major expenditure category. It appears that the mismeasurement of deliveries to final demand and value added in a few problem industries explains most of the broad movements in the aggregate discrepancy. In the following section, we discuss the metrics used to find an optimal combination of the GDP and GDI data to create an integrated data set. These metrics are based on standard economic arguments. We find that a mixture of data that do not aggregate either to GDP or to GDI appears to generate a data set that yields the best results. The fifth section concludes.

8.2   Methodology and Data

The main goal of the chapter is to construct and compare *consistent, integrated data sets* of the U.S. economy. We take "data set" to mean detailed information on the gross output, value added, final demand expenditures, and use of intermediate inputs by industry. We define a "consistent" data set to be one where the underlying components are based on the same definitions and industry classifications. And by "integrated," we mean that, despite the numerous data sources employed, the estimates conform to the accounting identities linking production, income, and expenditures.

Integration is not a unique transformation of the data, so different assumptions and methods to enforce integration can yield different estimates. We have built into our integration technique "tuning parameters" that summarize the specific assumptions that we use to obtain unique estimates. Adjusting these "tuning parameters" allows us to obtain different consistent, integrated data sets. In section 8.3, we compare two data sets based on polar assumptions: one integrates the data assuming that detailed GDP expenditures are correct; the other case assumes that income by industry (summing to GDI) is correct. In section 8.4, we estimate numerous data sets by varying the tuning parameters between the polar cases to compare their performance on predefined criteria.

It should be noted that the integration exercises are carried out on nominal data and that any comparisons made in real terms are based on the same deflators applied to either side of the comparison. Issues concerning how price and quantities can be consistently aggregated are considered in Moyer, Reinsdorf, and Yuskavage (chap. 7 in this volume).

The rest of this section describes the data and method employed to conduct our analysis. The first subsection illustrates our input-output system

that defines the components of our data set. The second subsection describes the sources of our initial estimates of these components and the manipulations we made to make them consistent. The final subsection describes the methodology used to integrate the source data to satisfy the constraints in our input-output system.

8.2.1   Our Input-Output System

The input-output system that describes the data set used in this study is shown in figure 8.2. Domestic industries, represented as the first N rows of the table, produce gross output (vector $\mathbf{Y}$) and deliver it to final demand (matrix $\mathbf{F}$) or to other domestic industries, (matrix $\mathbf{I}$), which use it as intermediate inputs in their production processes. The fact that the sum of each industry's deliveries to final demand and to other industries equals its gross output is called the *gross output identity.* The value added of an industry equals its gross output less the sum of its use of intermediate inputs (*value added identity*). The sum across industries of deliveries to final demand equals GDP (*GDP identity*), and the sum of value added across industries equals GDI (*GDI identity*). The *reconciliation identity* that integrates the system is that GDP equals GDI.

The first N rows of the system represent flows of goods from domestic industries. In order to simplify the exposition of our analysis, we account for the flows of imported goods in a nonstandard fashion: imported goods that are used in the production process of domestic industries or that are delivered to final domestic purchasers are the product of a separate industry,
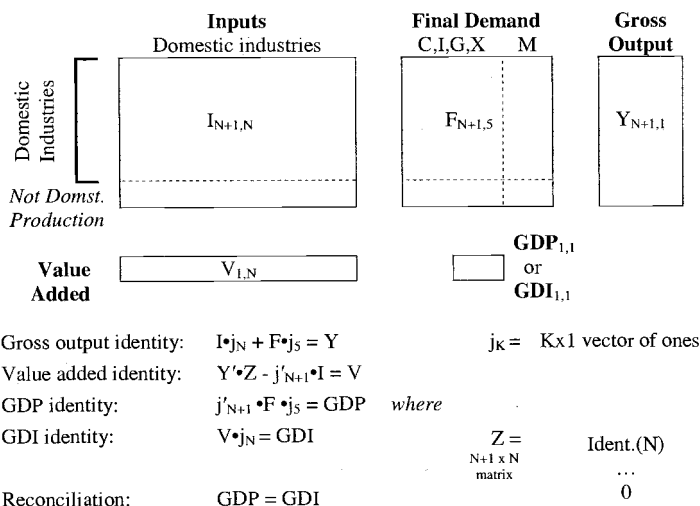

Gross output identity: $\mathbf{I} \bullet \mathbf{j}_N + \mathbf{F} \bullet \mathbf{j}_5 = \mathbf{Y}$ $\mathbf{j}_K = $ K×1 vector of ones

Value added identity: $\mathbf{Y}' \bullet \mathbf{Z} - \mathbf{j}'_{N+1} \bullet \mathbf{I} = \mathbf{V}$

GDP identity: $\mathbf{j}'_{N+1} \bullet \mathbf{F} \bullet \mathbf{j}_5 = \text{GDP}$ *where*

GDI identity: $\mathbf{V} \bullet \mathbf{j}_N = \text{GDI}$ $\mathbf{Z} = $ Ident.(N)
                                                     N+1 × N   ...
                                                      matrix   0

Reconciliation: $\text{GDP} = \text{GDI}$

Fig. 8.2   Input-output system

called *Not Domestic Production,* which is the last row of the upper blocks. Deliveries of imports to domestic industries or to domestic purchasers are positive entries in the input-output system. The final demand category, imports, has an offsetting negative entry, so that the gross output of imports is zero. Note that, by definition, domestic industries do not deliver any output to the final demand category imports, and so the first N rows of the import column contain zeros.

In addition, used and secondhand goods and scrap show up in the input-output accounts. They are used as intermediates to the production process and are either delivered to or supplied by the final demand categories. They do not represent new production, so, like imports, their gross output equals zero. Negative entries represent net suppliers of the goods; positive entries represent net users. For example, businesses scrap some of their equipment each year, so the final expenditure category, business fixed investment, is a net supplier of used and secondhand goods and scrap. These commodity flows are also included in the pseudoindustry Not Domestic Production.

8.2.2   Developing a Consistent Initial Data Set

In order to conduct our analysis, we need to populate the elements of the input-output system with initial values using consistent definitions. As described below, these initial values come from different published sources that do not match precisely in terms of definitions, accounting conventions, basis for data collection, or product and industry classifications. The GDP and GDI data for the years 1977 through 2001 come from the recently released benchmark National Income and Product Account (NIPA) data. Other data were adapted or created from the latest published data source from the BEA.

Value Added by Industry

Value added for farms, private households, and owner-occupied housing comes directly from the NIPAs. Value added for owner-occupied housing was subtracted out of the real estate industry and placed in its own industry (before further aggregation). For other industries, estimates of value added by industry are sourced from the BEA's 2003 GDP-by-industry data set. Pre-1987 data were concorded to the 1987 Standard Industrial Classification (SIC) as in Bartelsman and Beaulieu (2004). All of the income components were adjusted proportionately so that they sum to the latest aggregate estimates.

Value added in the real estate industry was also adjusted to exclude the imputed rental value of capital equipment and structures owned by nonprofit institutions. Instead, this imputed income was distributed to industries according to estimates of the compensation paid by nonprofit institutions by industry, as estimated in Bartelsman and Beaulieu (2004). Redistributing this income is useful because the final expenditures on many of the prod-

ucts produced by nonprofit institutions are not identified as to whether they were produced in the nonprofit sector or in the business sector, so these expenditures will not show up as coming from the real estate sector.

In putting together its GDP-by-industry data set, the BEA had to adjust some of its source data to put the data set consistently on an establishment basis. In particular, the original information on corporate profits, nonfarm proprietors' income, net interest paid, and capital consumption allowances is measured on a firm basis (U.S. Bureau of Economic Analysis 2001, pp. M21–M22). Other data, such as gross output and compensation paid, are measured on an establishment basis. The same income components collected on these two bases for the same industry will differ when firms in that industry have extensive operations in different lines of work. Data collected at the establishment level will split a multiestablishment firm into different industries, but data collected on a firm basis will put all of the firm's operations into one industry. For its GDP-by-industry data set (2003), the BEA adjusted the source data to put all of it on an establishment basis using a cross-classification table. But these are difficult adjustments to make, and this adjustment could be a source of error in allocating domestic data among industries. The finance industry is one where the distinction between firm and establishment data is particularly important (see Bartelsman and Beaulieu 2004).

Deliveries to Final Demand by Industry

No published data on deliveries to final demand by industry exist, so estimates based on detailed NIPA expenditure and input-output data had to be developed. First, detailed NIPA data on all expenditures, except software investment, construction, and inventory investment, were allocated to the input-output tables' commodity classification system. These mappings are called "bridge tables," the construction of which is described in detail below. The second step involves dividing final expenditures between domestically produced and imported commodities. Third, estimates of deliveries of commodities were converted to deliveries by industries. The domestic production of each commodity is converted to an industry basis using the 1987 and 1992 make tables, and these industries are then aggregated to the definitions in appendix table 8A.1. Imports of all commodities are aggregated into one industry, called Not Domestic Production.

The method used to estimate the bridge tables differs by expenditure category. For personal consumption and equipment investment (including residential equipment), detailed bridge tables were published by the BEA for 1987 and 1992. These bridge tables include the fraction of expenditures due to transportation and trade margins; these margins are treated as a separate commodity delivered to the specific expenditure category. For exports, imports, and government expenditures, bridge tables were created by assigning commodities to specific NIPA categories using the 1987 and 1992

use tables to estimate specific proportions. For exports and imports of goods, NIPA expenditures were disaggregated to more detailed census categories using information in the Bureau of the Census report on international trade in goods and services; input-output (I-O) commodities were assigned to these more detailed census categories. Export margins for wholesale trade and goods transportation were allocated across expenditure categories in the same proportion as total margins to all goods exports as shown in the use tables.

Bridge tables for government consumption were built by first assigning the consumption of fixed capital and the compensation paid to general government employees, excluding own-account investment to the general government industry. Compensation paid to employees for own-account investment is treated with other government investment. Commodities with positive values in the I-O use tables were assigned to government purchases of intermediate durables, nondurables, and services, depending on the commodity's characteristics. Commodities with negative values in the I-O use table were assigned to government sales.[3] Netting out government sales from intermediate purchases yields government consumption excluding its own value added. The NIPA data on federal nondefense, nondurable consumption were augmented with data from the Energy Information Agency to account for purchases and sales from the Strategic Petroleum Reserve. As with trade, margins were distributed to all expenditure categories in fixed proportions.

Bridge tables for government investment were created by first splitting own-account investment into equipment and structures using pre-revision data on compensation paid to force-account construction. Own-account investment originates from the general government. The remaining investment in structures was assigned to the construction industry, and the remaining investment in equipment was split among commodities using relative proportions in the 1987 and 1992 I-O use tables.

Imports are different from other expenditure categories in that all imports are counted as coming from one industry. However, it is necessary to allocate a fraction of imports to the domestic final purchases categories and the rest to intermediate inputs to domestic production in order to estimate the fraction of each commodity delivered to final demand that was produced domestically versus imported. This split was done by assuming that the fraction of an imported commodity delivered to final demand categories versus to domestic industries is the same as that observed in the I-O use tables. The rest of final demand is then assumed to be produced domestically.

The production of each commodity was then converted to an industry basis using the 1987 and 1992 I-O make tables. We assumed that the pro-

---

3. The NIPAs provide more detail on intermediate purchases for federal defense and the sales by state and local governments that are used to refine these assignments.

portion of each commodity that was produced by the I-O industry was the same as indicated in the make tables. Using 1987 and 1992 data produces two estimates. For the years 1987 and before we used the estimates based on the 1987 tables; for the years 1992 and after we used the estimates based on the 1992 tables. For the years in between, we used a weighted average of the two, where the weights are based on the distance from each benchmark year. These industry estimates were then aggregated to the industry definitions as in appendix table 8A.1.

Residential and nonresidential investment in structures by industry had to be estimated in a different manner than would follow from the published I-O tables. Some expenditure categories were assigned directly to specific industries: drilling and exploration to mining, mobile homes to the appropriate manufacturing industry, and commissions to real estate.

The I-O tables appear to suggest that the remainder of investment in structures originates in the construction industry, but this is not correct. For construction, the I-O tables make an exception to the rule that production is classified according to the primary output of an establishment. Instead, the tables classify all construction regardless of the primary output of an establishment to the construction industry, a classification scheme known as activity based. Most of the rest of the input-output data are essentially organized on an establishment basis.[4] Figure 8.3 illustrates the problem with mixing establishment-based classifications and activity-based classifications: domestic investment in structures, excluding government own-account investment in structures exceeds the BEA's estimate of gross output in the construction industry. Consequently, we have to estimate how much of private structures investment originates in the construction industry versus other industries.

The value of deliveries to final demand by the construction industry was calculated as a fraction of the BEA's estimate of gross output. This equals the interpolated values of one minus the ratio of receipts for maintenance and repair to total sales in the Censuses of Construction (1977, 1982, 1987, 1992, and 1997).

The remainder of investment in structures was assigned to other industries based on their share of employment of construction workers in 2001 (from the BLS occupational survey) times the BEA's estimate of the real wealth stock of structures by industry. Including the real wealth stock allows the indicators used to allocate the estimate of force-account construction to vary over time.

Software investment was allocated across industries by first splitting in-

4. Farms and real estate services are the other industries in the input-output tables that are defined on an activity basis instead of an establishment basis. The farm industry, however, is consistently treated in the NIPAs. All royalty income, regardless of its origination, is counted in the real estate industry, but this is the same treatment in the GDP-by-industry data. Thus, adjustments are not necessary to improve the consistency of these industry estimates.
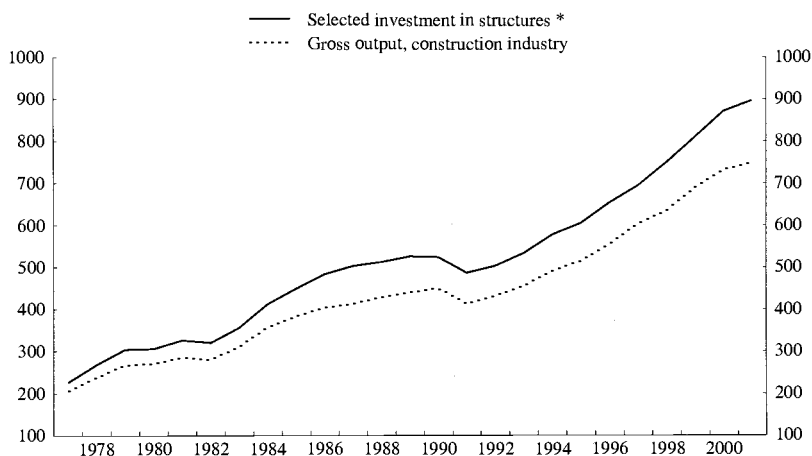
Fig. 8.3   Measures of construction activity (billions of dollars)

Note: Selected investment in structures excludes investment in mining and exploration, manufactured homes, commissions, and government own-account investment.

vestment into two components: own account and purchased software using the BEA's detailed new investment-by-industry data. Own-account investment was then allocated across industries using these data. Purchased software was distributed to industries using the 1987 and 1992 make tables; 98 percent of the production of purchased software in 1992 was assigned to the data-processing services industry, SIC 737.

Inventory investment was allocated to industries based on published NIPA data. Farm inventories were assigned to farms. Manufacturing inventory investment was allocated among manufacturing industries using book value data from the Annual Survey of Manufacturers (ASM). ASM data reported on a NAICS basis or on the 1977 SIC were concorded to the 1987 SIC using available concordances. Wholesale and retail trade inventories were simply assigned to the trade industry. The remainder of inventory investment was allocated among other industries using data from the Sources of Income (Department of Treasury) for 1995–97. Shares for other years were assumed to equal either the 1995 or 1997 value.

Table 8.1 describes how well our bridge tables translate the available detailed NIPA expenditure data into deliveries to final demand by industry. As shown in the first row, personal consumption expenditures (PCEs) were $3,100.2 billion in 1987. The BEA breaks up total PCE into 141 categories, such as sporting equipment, sugar and sweets purchased for off-premise consumption, and spending on theater and opera performances. On average, each of the 141 detailed categories was divided among 6.1 industries. One quarter of PCE was in expenditure categories that were allocated all to just one industry. Another 22.5 percent of PCE was in categories where over

Table 8.1   **Splitting NIPA expenditure categories (1987) across 26 industries**

	Billions of 1987 $	Number of expenditures categories	Industries per expenditure category	Dominant industry's share of expenditure category						Standard Deviation[a]
				100%	95–100%	85–95%	75–85%	50–75%	Under 50%	
Personal Consumption	3,100.2	141	6.1	24.9	22.5	5.2	8.1	28.7	10.6	2.4
Equipment Investment	326.7	27	11.5	8.8	0.0	16.0	51.8	17.9	5.5	6.7
Software Investment	29.0	1	23.0	0.0	0.0	0.0	0.0	100.0	0.0	11.1
Structures Investment	402.1	5	5.8	3.3	1.3	6.1	0.0	89.3	0.0	10.0
General Government	591.2	1	1.0	100.0	0.0	0.0	0.0	0.0	0.0	0.0
Other Government	408.4	39	7.2	18.8	15.8	0.7	5.0	12.1	47.7	5.6
Exports, Goods	257.5	124[b]	8.3	0.0	8.6	33.9	18.8	27.3	11.4	2.4
Exports, Services	106.4	7	5.6	29.8	18.0	9.6	0.0	42.6	0.0	7.1
Imports, Goods	–414.8	128[b]	5.5	58.1	27.2	2.3	3.0	9.5	0.0	3.0
Imports, Services	–94.4	7	4.4	48.9	12.2	0.0	0.0	20.1	18.8	9.9
Inventory Investment[c]	27.1	4	6.8	53.8	0.0	0.0	18.2	0.0	0.0	57.6
GDP	4,739.5	484	6.9	26.4	14.5	7.0	9.9	30.4	11.8	3.3

[a] Average standard deviation of deliveries by industry to these major expenditure categories assuming a 10 percent standard deviation in the bridge tables. Deliveries by industry are subsequently calculated by controlling the bridge tables to the actual values of the detailed expenditure categories.

[b] NIPA expenditure categories are divided into more categories using data from the monthly report on international trade in goods and services.

[c] Dominant industry's share calculated using absolute values instead of actual values. Standard deviation calculated using normal, additive errors.

95 percent of the category was allocated to one industry (fifth column). Only 10.6 percent of PCE was in categories that were so diffuse that the largest industry did not account for half of the category (ninth column).

The bridge tables contain a lot of structure that constrains how relative errors in the bridge tables can affect our estimates of deliveries to final demand by industry. For example, the value added of the general government, which the BEA publishes, maps to only one industry, and so, conditional on this published value, this category cannot contribute to an error in our estimates. To take another example, PCE radio and television repair services are estimated to be produced by three industries—personal services, business services, and machinery manufacturing—with personal services accounting for 95 percent of final demand. As a result, for this category of consumption, a large relative error in the bridge table for business services and machinery manufacturing can have only a small effect on the estimated deliveries of personal services.

To see how errors in the bridge table can translate into variation in our estimates of deliveries to final demand by industry, we performed the following experiment. We multiplied the cell values in our 1987 bridge tables by lognormally distributed errors so that the standard deviation of the cell values was 10 percent, and then we recontrolled the bridge tables so that the sum across industries equaled the published values of the detailed expenditure categories. We then recalculated the implied deliveries to final demand. We repeated this procedure 2,500 times. As shown in the last column of the table, a 10 percent random error in the bridge tables translates to only an average variation of deliveries by industry to PCE of 2.4 percent.

Other major categories are not measured as well. For equipment investment a 10 percent error in the bridge table leads to an average standard deviation of 6.7 percent in deliveries to final demand by industry. This weaker performance is probably due to the poorer precision in the equipment investment bridge table. On average, there are 11.5 industries per category, and three-quarters of equipment investment is spread among categories where the dominant industry accounts for less than 75 percent spending. For all of GDP, a 10 percent error in the bridge tables translates to a 3.3 percent error in deliveries to final demand by industry.

Gross Output by Industry

Estimates of gross output by industry come mainly from the published GDP-by-industry data, except for farms, owner-occupied housing, general government, and households, which are available or easily estimated from NIPA data. In a few early years, the estimate of value added by the legal services industry was higher than the estimate of gross output. To allow our analysis to proceed, we boosted the value of gross output so that it exceeds value added by at least 5 percent, a figure consistent with the 1987 I-O use table.

Intermediate Inputs

The starting point for constructing the intermediate block of the consistent data set is the use table from the published BEA benchmark I-O data. Unlike the vectors and matrices for gross output, deliveries to final demand, and value added, the initial values for the intermediate block, **I**, are calculated only for the base years 1982, 1987, and 1992. Initial values for other years are developed iteratively using results from the balancing routine described in the next subsection.

Initial values for the base years were calculated twice and then averaged to get one estimate. The first estimate allocates the vector of gross output less deliveries to final demand ($\mathbf{Y} - \mathbf{F}$) across the columns of **I** in proportion to the values observed in the 1982, 1987, or 1992 use tables. The second estimate allocates the vector of gross output less value added ($\mathbf{Y}' - \mathbf{V}$) across the rows of **I**, also in proportion to the values observed in the corresponding use tables. These two estimates, one of which can be thought of as consistent with the expenditure-side data, the other as consistent with the income-side data, are then combined by taking a geometric average of the two values cell by cell.

The resulting benchmark-year initial estimates of **I** are adjusted to subtract out the intermediate value of software purchases, which are now counted as final demand (see Bartelsman and Beaulieu 2004), and adjusted to allocate own-account construction to the appropriate industries. Further, the values in the columns from the use table for transportation margins and distribution margins are entered as intermediate purchases by the industry purchasing the relevant input and as sales to other industries by the "margin industries," such as water and rail transport or retail trade.

8.2.3   Integrating the Data

The consistent I-O data set populated with initial values is adjusted, or integrated, so that the various constraints in the input-output system are satisfied with cell values "close" to the initial estimates. Specifically, we choose values for each element in the I-O system to minimize the weighted sum of squares of the difference with its initial estimate subject to the linear constraints. The inverse of the weights equals the absolute value of the cell times a "tuning" parameter; these tuning parameters are what we use to control the integration process. The closer the tuning parameter is to zero, the more we restrict the final estimate to lie close to the initial estimate. If the tuning parameter equals zero, the value of the cell is not adjusted. This solution technique is a straightforward generalization of the least squares method first proposed by Stone, Champernowne, and Meade (1942).

Formally, denote the initial estimates of each element of the vectors and matrices of the I-O system with a bar. We solve

$$(1) \qquad \min_{\{\mathbf{Y}_t, \mathbf{F}_t, \mathbf{V}_t, \mathbf{I}_t\}} \sum_{i=1}^{I} \frac{1}{\sigma^{\mathbf{Y}} |\overline{\mathbf{Y}}_{it}|} (\mathbf{Y}_{it} - \overline{\mathbf{Y}}_{it})^2 + \sum_{i=1}^{I} \frac{1}{\sigma^{\mathbf{F}} |\overline{\mathbf{F}}_{it}|} (\mathbf{F}_{it} - \overline{\mathbf{F}}_{it})^2$$

$$+ \sum_{i=1}^{I} \frac{1}{\sigma^{\mathbf{V}} |\overline{\mathbf{V}}_{it}|} (\mathbf{V}_{it} - \overline{\mathbf{V}}_{it})^2$$

$$+ \sum_{i=1}^{I} \sum_{j=1}^{I} \frac{1}{\sigma^{\mathbf{I}} |\overline{\mathbf{I}}_{ijt}|} (\mathbf{I}_{ijt} - \overline{\mathbf{I}}_{ijt})^2$$

$$\text{s.t.} \quad \mathbf{Y}_{it} = \mathbf{F}_{it} + \sum_{j=1}^{I} \mathbf{I}_{ijt}$$

$$\mathbf{Y}_{jt} = \mathbf{V}_{jt} + \sum_{i=1}^{I} \mathbf{I}_{ijt},$$

where σ denotes preset tuning parameters, i denotes row, j denotes column, and t denotes time. If σ equals zero, then the weight becomes a Lagrange multiplier and the fact that the cell value equals its initial value becomes another restriction in the minimization problem.

As indicated in equation (1), because the inverse of the weights are proportional to the initial values, initial values that are equal to zero are restricted to remain zero. In our application we restrict the values of σ to be the same for all elements of the same vector or matrix. For example, all values of σ for the value-added vector are equal to σ^V, with one exception, which is described in the next section. One could also allow these parameters to differ across industries, for instance, if there was some idea that some industries were measured better than others, but we do not pursue this angle. Finally, it should be obvious from equation (1) that only the relative values of σ matter; doubling all of them does not change the solution. Thus, we standardize the parameters by setting $\sigma^I = 1$. Furthermore, to focus our analysis we only consider $\sigma^Y = 0$; this leaves a pair of tuning parameters $\{\sigma^F, \sigma^V\}$ to vary.

Other solution techniques have been used for similar problems. In particular, a popular routine is the so-called RAS iterative solution. Although Leontief in 1941 had already suggested a biproportional form for the relationship between the values taken by an I-O matrix at different points of time it became popular after it was used in R. A. Stone's Cambridge Growth Project in the early 1960s. It then became familiar under the name "RAS technique." In the traditional RAS or biproportional balancing method used for integration, differences between "control" totals and the sum of unadjusted data in one dimension are iteratively applied to proportionally adjust data in the other dimension until both restrictions are satisfied within a prescribed tolerance level.

The starting point of the algorithm is a given matrix **A** with semipositive rows and columns and strictly positive vectors **u** and **v**. One first multiplies each row by a scalar that will make the row sum equal the row constraint. Next one multiplies each column of the resulting matrix $\mathbf{A}^1$ by a scalar that will make its sum equal its constraint. This gives matrix $\mathbf{A}^2$, which serves as starting point for the next iteration. In general the process can be described as follows:

$$\mathbf{A}^{2t+1} = \hat{\mathbf{r}}^{t+1}\mathbf{A}^{2t}$$

$$\mathbf{A}^{2t+2} = \mathbf{A}^{2t+1}\hat{\mathbf{s}}^{t+1} = \hat{\mathbf{r}}^{t+1}\mathbf{A}^{2t}\hat{\mathbf{s}}^{t+1}$$

$$\hat{\mathbf{r}}^{t+1} = \frac{\mathbf{u}_i}{\sum_{j=1}^{n} a_{ij}^{2t}}$$

$$\hat{\mathbf{s}}^{t+1} = \frac{\mathbf{v}_j}{\sum_{i=1}^{m} a_{ij}^{2t+1}},$$

where $\hat{\mathbf{r}}$ and $\hat{\mathbf{s}}$ indicate a matrix with diagonal elements $\mathbf{r}_i$ such that $\mathbf{r}_i = \mathbf{u}_i / \Sigma\, a_{ij}$.

Bacharach (1970) established the existence of a solution to the RAS process under weak conditions and the uniqueness of the solution unconditionally. Also, for the first time stochastic elements (specification errors) were introduced. It is further shown that this RAS method is equivalent to a minimization problem with as solution the biproportional estimates r and s.

Bishop, Fienberg, and Holland (1975) present a methodology that is developed specifically for the analysis of categorical (qualitative) data in multidimensional contingency tables. Although they focus on maximum likelihood estimation and log-linear models, elements of the development and statistical properties of the RAS method are discussed. They discuss both complete and incomplete tables (including iterative proportional fitting of log-linear models) and pay special attention to marginal homogeneity and symmetry.

Unlike our technique, the iterative RAS method does not have a natural role for tuning parameters.[5] In addition, a problem with the RAS method arises when the controls do not sum to the same total; in practice, one or both of the controls are adjusted to coincide before the RAS procedure is

---

5. The iterative RAS solution is the solution of a minimization problem subject to the biproportional constraint, where instead of minimizing quadratic differences, the entropy kernel is used. Schneider and Zenios (1990) credit a Russian mathematician Bregman for this result, although the fact that the first-order conditions for the minimization problem yield the RAS iterative solution is not difficult to illustrate; see, for example, Günlük-Şenesen and Bates (1988). One could therefore weight the entropy kernel to allow for tuning parameters, though this would complicate the iterative technique to arrive at a solution. Bartelsman and Beaulieu (2003) explore some of the implications of the choice of balancing technique; see also Schneider and Zenios.

applied. In our method, the "controls" are not adjusted before minimization; instead, our routine adjusts the controls simultaneously with the other estimates as specified by the tuning parameters.
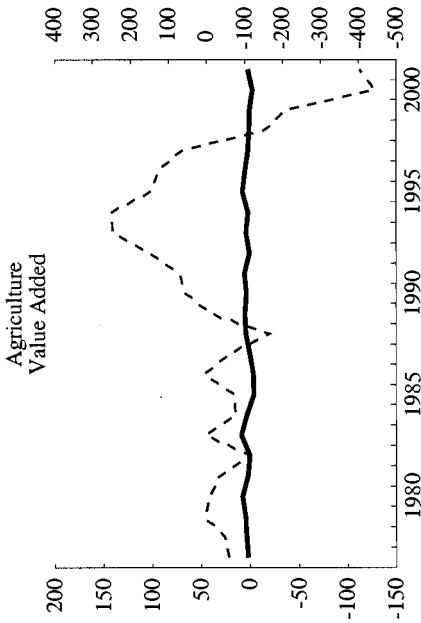
As noted in the previous subsection, our estimation procedure is dynamic in that our initial estimates of $\bar{\mathbf{I}}_t$ depend on the final results for other years when $t \neq 1982, 1987, 1992$. We first estimate the system for 1982 and then move backward in time to 1977 and forward in time to 1986, using the final estimate of $\mathbf{I}_{t\pm1}$ as a basis for $\bar{\mathbf{I}}_t$. Specifically, $\bar{\mathbf{I}}_t$ is calculated by adapting $\mathbf{I}_{t\pm1}$ for demand changes in the various columns by multiplying each cell of $\mathbf{I}_{t\pm1}$ by the ratio of real gross output of column j in period t to real output of j in period $t \pm 1$. The matrix $\mathbf{I}_{t\pm1}$ is also adapted for price changes in the various rows by multiplying each cell by the ratio of the gross output deflators for row i in period t to the output deflator in period $t \pm 1$. The same process is repeated starting in 1987 for the years 1983–91 and starting in 1992 for the years 1988–2001. This produces two sets of estimates, in current dollars, for 1983–86 and for 1988–91; these estimates are averaged to obtain one series of $\bar{\mathbf{I}}_t$ for 1977–2001.

8.3   Results Controlled to the GDP or GDI Data

Equation (1) was first estimated under two sets of tuning parameters. The first set, $\{\sigma^{\mathbf{F}} = 0; \sigma^{\mathbf{V}} = 1\}$,[6] means that we controlled the estimates to the expenditure-side data, and it leads to estimates of industry value added and deliveries to final demand that add to GDP. We allow the initial income-side value-added estimates to inform our final estimates, but with $\sigma^{\mathbf{V}} = \sigma^{\mathbf{I}} = 1$ the routine treats the estimates of value added symmetrically with the initial estimates of $\mathbf{I}$. The second set of tuning parameters, $\{\sigma^{\mathbf{F}} = 1; \sigma^{\mathbf{V}} = 0\}$, implies that we controlled the estimates to the income-side data; it leads to estimates of industry value added and deliveries to final demand that sum to GDI. In both cases, $\sigma^{\mathbf{Y}}_{\text{MISC}} = \sigma^{\mathbf{V}}_{\text{MISC}} = 0$ because the income and gross output of these industries are already integrated between the expenditure and income accounts. Early experiments with the estimation procedure gave estimates for the Not Domestic Production industry that tended to drift. With both negative and positive values for deliveries of this series tied down only to sum to zero, the estimates of this industry can be volatile. As a result, $\sigma^{\mathbf{V}}_{\text{NDP}} = \sigma^{\mathbf{F}}_{\text{NDP}} = 0.00001$ if it otherwise is not equal to zero. Thus, we allow only small differences from the initial estimates for this industry, and it means that our estimate for the statistical discrepancy for imports essentially equals zero.

Figure 8.4 plots the difference of the two estimates for each industry's deliveries to final demand in the left-hand panels and the difference of the two estimates for each industry's value added in the right-hand panels. Using

6. Recall that in all of our estimates $\sigma^{\mathbf{Y}} = 0$ and $\sigma^{\mathbf{I}} = 1$.

Agriculture
Deliveries to Final Demand

Agriculture
Value Added

Mining
Deliveries to Final Demand
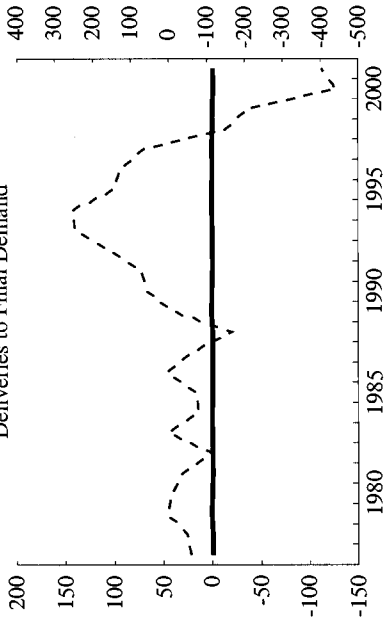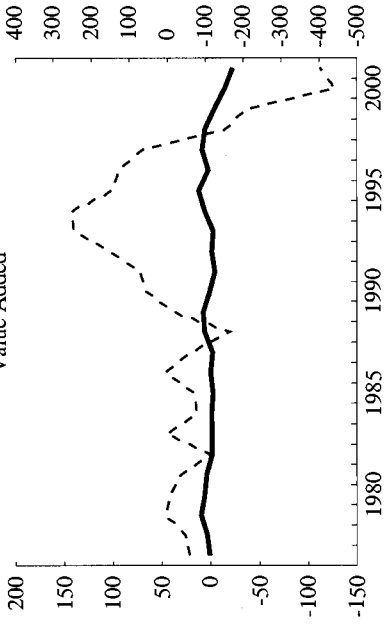
Mining
Value Added
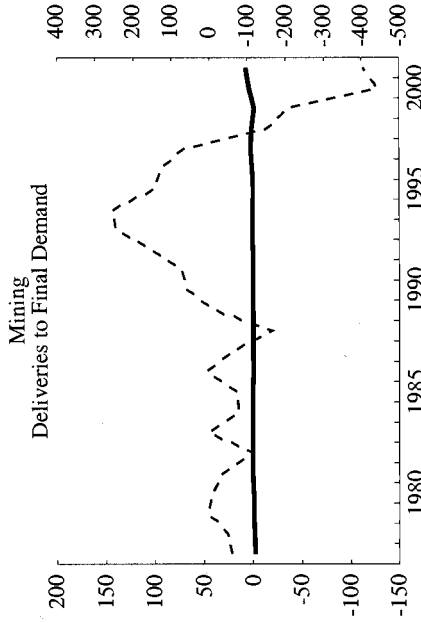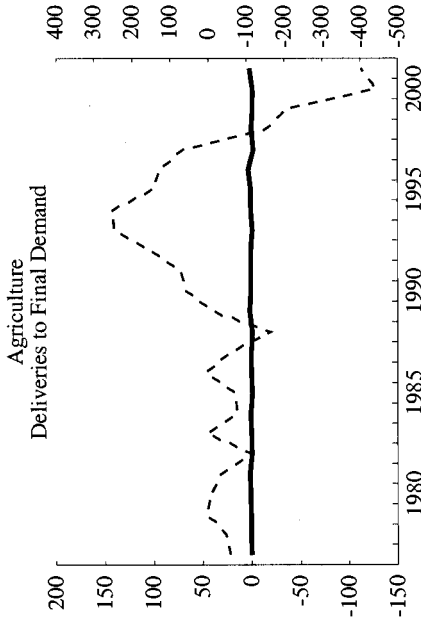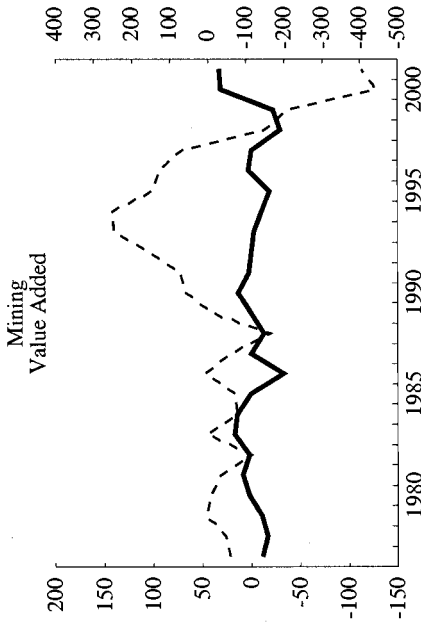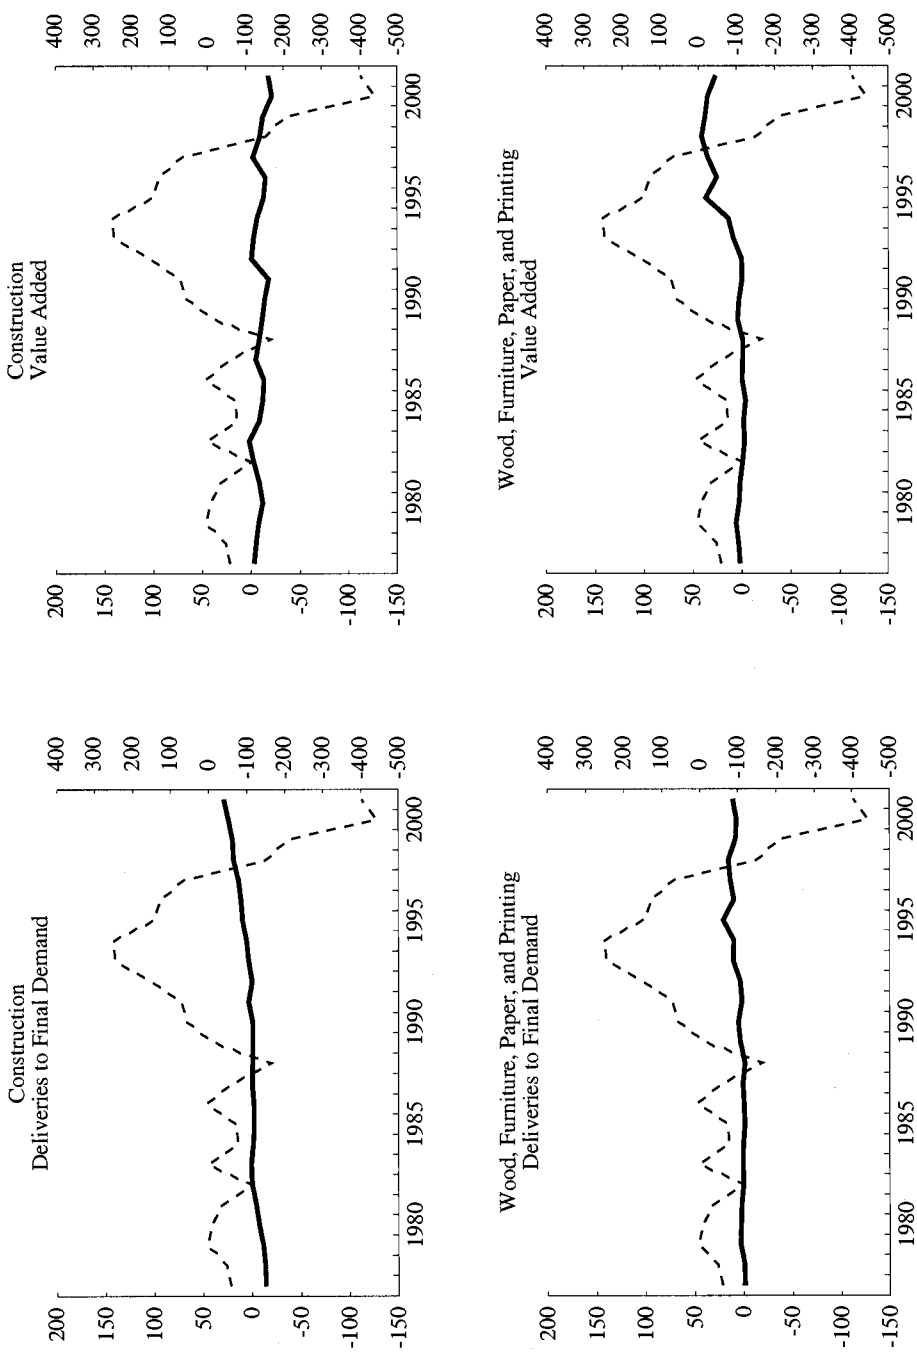
Fig. 8.4 Statistical discrepancy by industry (billions of dollars)

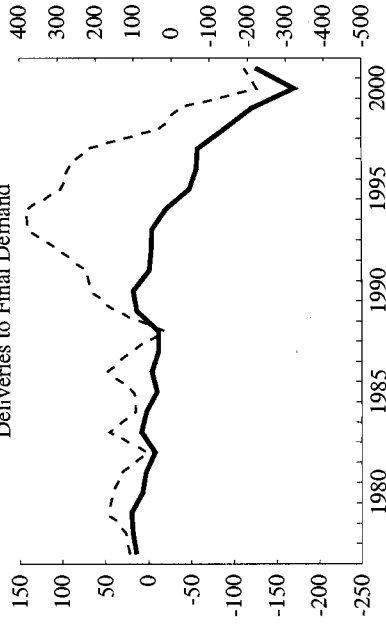*Note*: Industry discrepancy is shown by the solid line; aggregate discrepancy is shown by the dashed line.
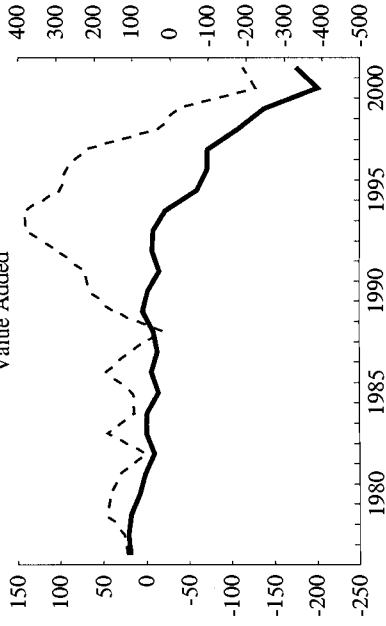
Primary Durable Mfg.
Value Added

Primary Durable Mfg.
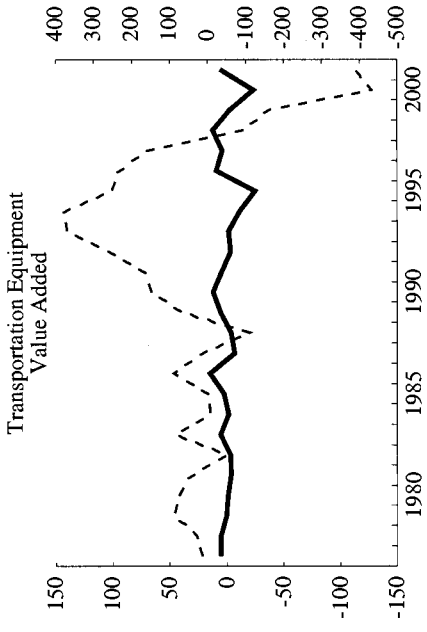Deliveries to Final Demand

Machinery and Instruments
Value Added

Machinery and Instruments
Deliveries to Final Demand

Transportation Equipment
Value Added

Transportation Equipment
Deliveries to Final Demand

Food and Tobacco
Value Added

Food and Tobacco
Deliveries to Final Demand

Fig. 8.4 (cont.)

Textiles, Apparel, and Leather
Value Added

Textiles, Apparel, and Leather
Deliveries to Final Demand

Chemicals, Refining, and Rubber & Plastics
Value Added

Chemicals, Refining, and Rubber & Plastics
Deliveries to Final Demand

Transportation
Deliveries to Final Demand

Transportation
Value Added

Communications
Deliveries to Final Demand

Communications
Value Added

Fig. 8.4   **(cont.)**

Utilities
Deliveries to Final Demand

Utilities
Value Added

Trade
Deliveries to Final Demand

Trade
Value Added

Finance and Insurance
Value Added

Real Estate
Value Added

Finance and Insurance
Deliveries to Final Demand

Real Estate
Deliveries to Final Demand

Fig. 8.4 (cont.)

Hotels and Other Lodging
Deliveries to Final Demand

Hotels and Other Lodging
Value Added

Personal Services
Deliveries to Final Demand

Personal Services
Value Added

Fig. 8.4 (cont.)

Health Services
Deliveries to Final Demand

Health Services
Value Added

Legal Services
Deliveries to Final Demand

Legal Services
Value Added

Fig. 8.4 **(cont.)**

Fig. 8.4 (cont.)

the convention used in the definition of the overall discrepancy, the figure plots the difference in the first measure, which aggregates to GDP, less the second measure, which aggregates to GDI. Each of these differences can be considered statistical discrepancies by industry. The economywide statistical discrepancy is also plotted in all of the panels.

For most industries, the industry discrepancies are small relative to the overall discrepancy. Three industries, however, stand out: Machinery and Instruments, Trade, and Finance and Insurance, where the pattern of deliveries to final demand and value added appear to move with the total discrepancy. Indeed, as shown in figure 8.5, the difference in value added of the combination of these three "problem" industries, moved up in the early 1990s and dropped sharply subsequently, more so than the total discrepancy. The coincidence with the discrepancy in deliveries to final demand is not as sharp. The difference in deliveries to final demand of the problem industries remained flat in the first half of the 1990s, but, like value added, the difference dropped sharply after 1996.

The fact that these three industries—Machinery and Instruments, Trade, and Finance and Insurance—show up as problem industries is not surprising. The Machinery and Instruments industry has evolved significantly over the last twenty-five years as productivity growth in high-tech industries has been substantial. Profit swings have been significant, and the adjustment of industry profits from a firm basis to an establishment basis is probably difficult. The semiconductor industry is particularly challenging as several firms have become "fabless." These firms develop products but contract o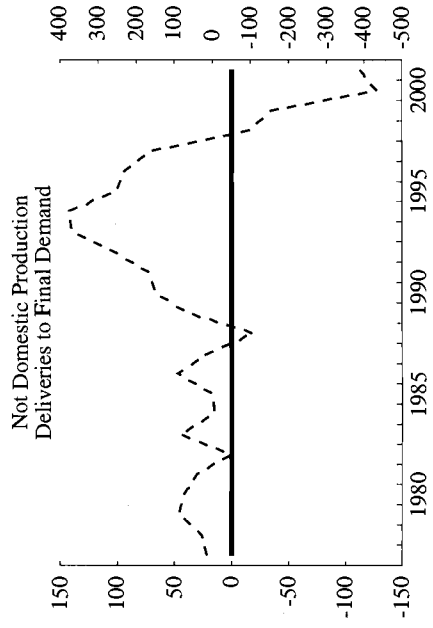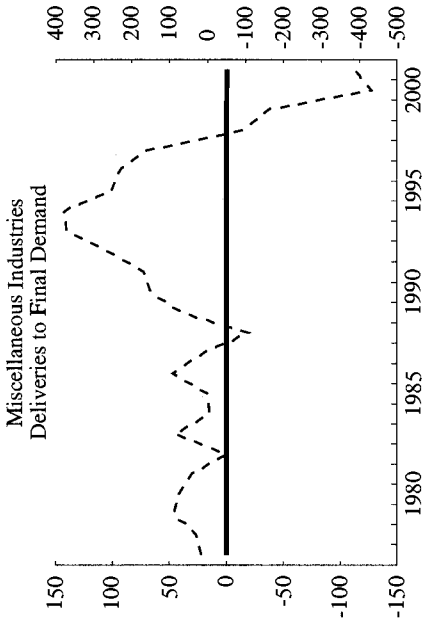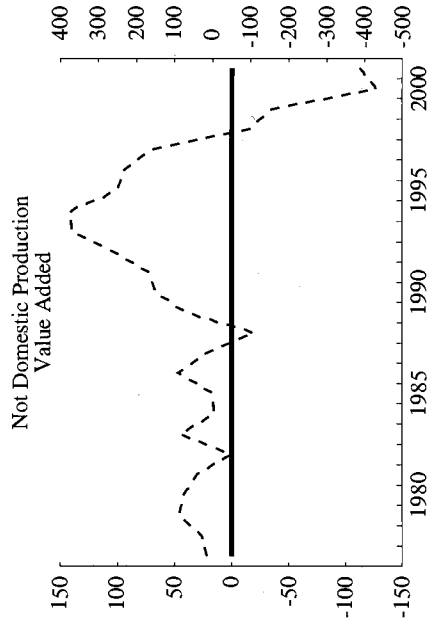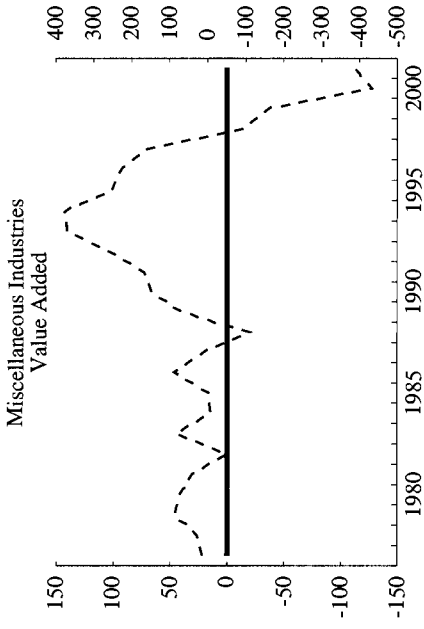ut their production to overseas fabrication plants. Morgan Stanley estimates that about 15 percent of the industry's worldwide revenue is derived from products outsourced to different firms (Edelstone et al. 2003); much of this figure represents U.S. firms contracting with overseas foundries. Morgan Stanley expects this share to double by 2010.

The difficulties with the Trade industry likely relate to the accounting for margins on products sold. To the extent that these differences represent margins on domestic products, there is a corresponding offset in the difference between the two measures in the domestic industries producing the output. If this is the reason for the discrepancy in the trade sector, then it cannot be a source for the economywide discrepancy. On the other hand, if the differences arise from different margins on imported products, difficulties in tracing these products from imports to deliveries to domestic purchasers could be a source of the overall discrepancy.

Finance and Insurance is clearly an industry fraught with measurement difficulties. A good deal of banking services is not explicitly charged for. Banks offer services like "free checking" to its customers because it can make money by lending the balances that customers leave in their accounts; customers choose to deposit their money in banks instead of lending it at higher rates to take advantage of the convenience of checking. The BEA has

Deliveries to Final Demand

- - - Total Discrepancy
—— Problem Industries *
······ Other Industries

Value Added

- - - Total Discrepancy
—— Problem Industries *
······ Other Industries

* Problem industries are: Machinery and Instruments, Trade, and Finance and Insurance.
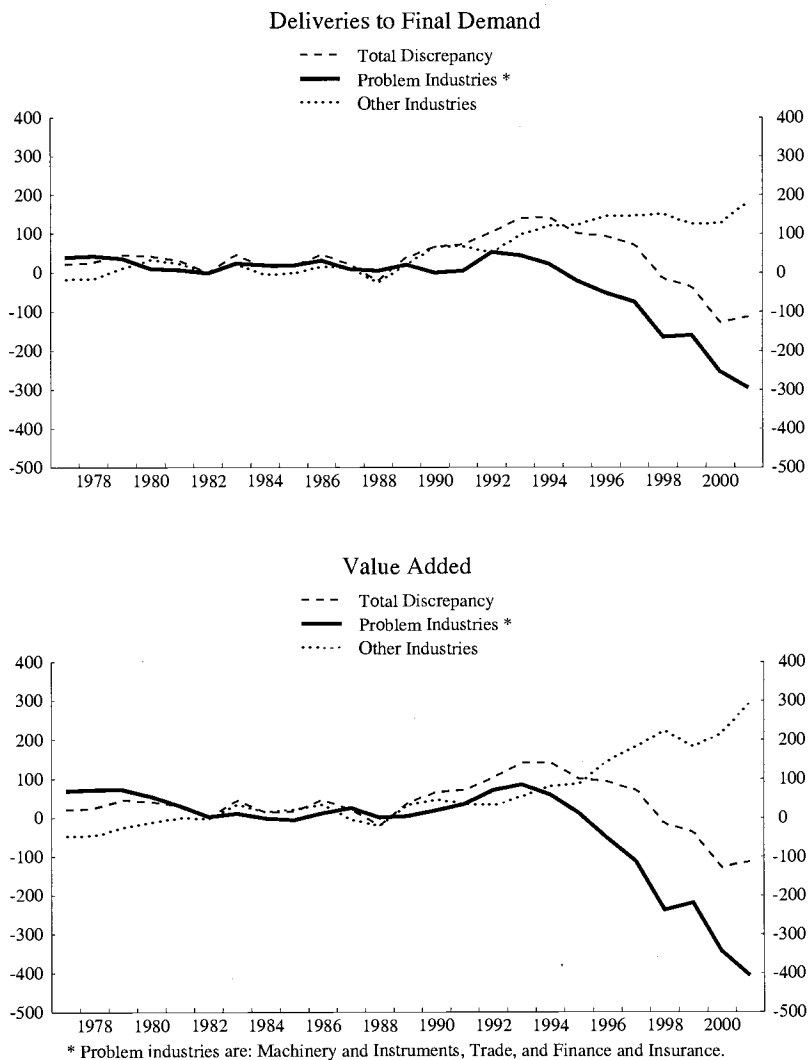
Fig. 8.5   Statistical discrepancy of problem industries (billions of dollars)
Note: Problem industries are Machinery and Instruments, Trade, and Finance and Insurance.

made substantial improvements to its estimation of imputed bank service charges in PCE and government consumption to account for these services (Fixler, Reinsdorf, and Smith 2003); however, the division of these services between final demand and intermediate inputs to business is probably still imprecise. The accounting for insurance services is likewise difficult. The same issue of imputed intermediation services arises in insurance. More-

over, the true value of insurance services is not realized only when claims are paid; there is a continual flow of services. Over the long run, the difference in premiums received less claims paid equals the services provided. How to estimate the evolution of these services over time is a thorny problem; the BEA has also improved its measures of deliveries to final demand of property-casualty insurance in the latest revision (Chen and Fixler 2003). On the income side, adjusting for capital gains has to be more difficult in the Finance and Insurance industry than in any other. Another complication may be the allocation of profits of large firms, such as GE, General Motors, and Ford, with establishments that operate in finance and in manufacturing.

A few other industries show some important differences that are not related to the overall discrepancy. Since 1995, deliveries to final demand of Chemicals, Refining, and Rubber and Plastics controlled to expenditure-side aggregates has risen sharply relative to estimated deliveries controlled to income-side aggregates, while for Communications the opposite is true. Over the same period, the value-added statistical discrepancy in Mining and in Health Services has increased rapidly, helping to offset some of the sharp decline in the statistical discrepancy of the problem industries.

Figure 8.6 plots the difference in the estimates of total deliveries to final demand by major expenditure categories. As is evident in the chart, essentially all of the run-up in the aggregate discrepancy in the first half of the 1990s occurred in PCE; much of the subsequent decline in the aggregate is also reflected in PCE. At the same time, however, the statistical discrepancy in private fixed investment also has trended down because of problems in the Machinery and Instruments industry. In 2001, there is an anomalous jump in the discrepancy in private fixed investment. Most of this is also in the Machinery and Instruments industry, but about $10 billion of this jump comes from Business Services, which includes software makers. As such, the post-Y2K slowdown in high-tech shows up more strongly in the data set controlled to income measures than in the data set controlled to expenditure measures.

8.4   Optimal Combination of the Data

In contrast to the exercise in the previous section, we now consider tuning parameters chosen to allow both value added by industry and final demand data to deviate from their initial estimates. The exercise is to search for a set of tuning parameters that provides an optimal result with respect to metrics based on desirable economic properties. The economic properties that we consider concern

- the equalization of returns to capital,
- the orthogonality of total factor productivity shocks, and
- the stability of the intermediate block.
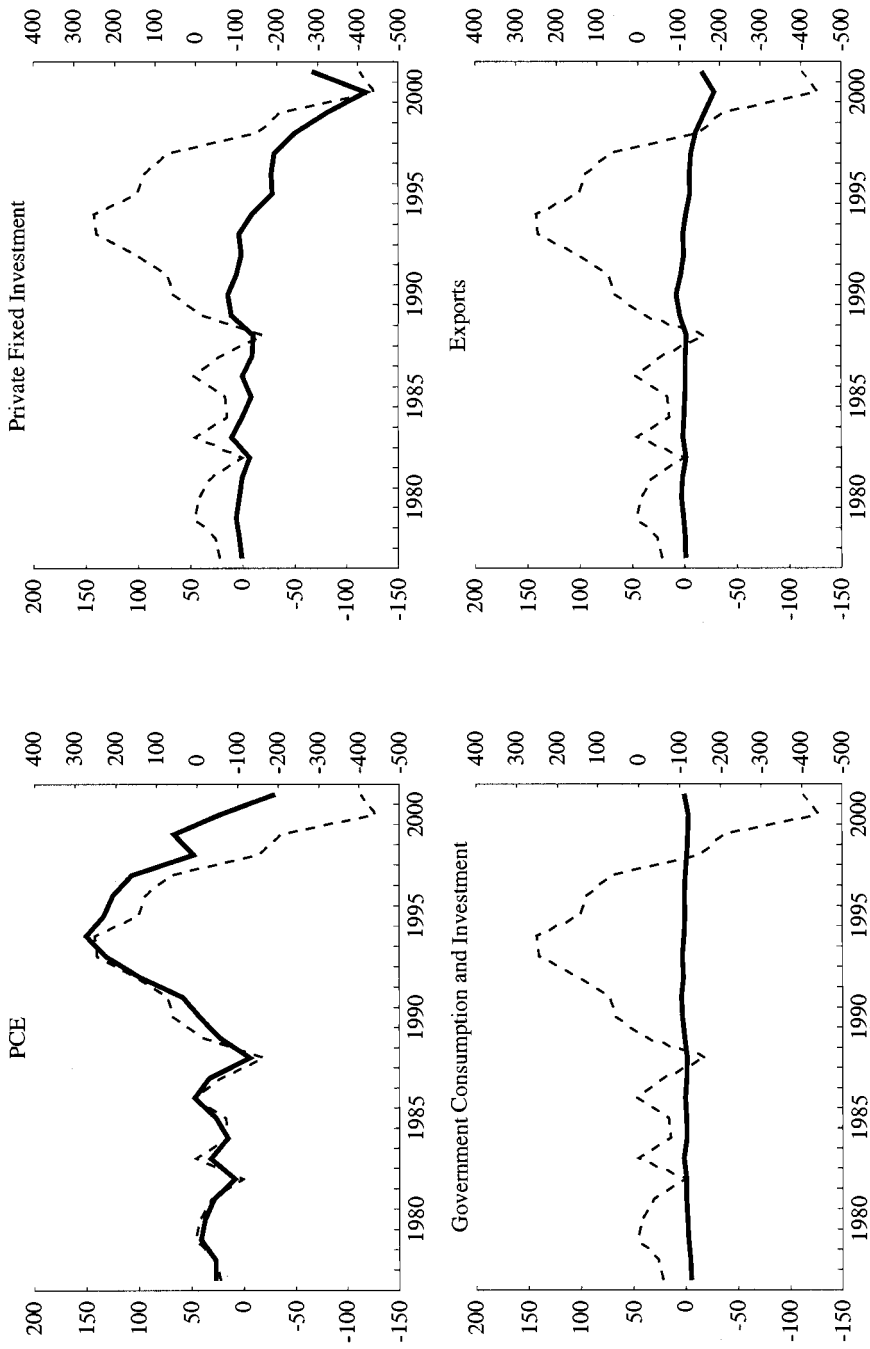
Fig. 8.6 Statistical discrepancy by expenditure category (billions of dollars)

Notes: Expenditure category discrepancy is shown by the solid line; aggregate discrepancy is shown by the dashed line. Differences in inventories are less than $1 billion in any year; differences in imports are restricted to be essentially zero.

*Differences in Inventories are less than $1 billion in any year; differences in imports are restricted to be essentially zero.

Our strategy is to estimate a series of consistent, integrated data sets under different assumptions for the tuning parameters $\{\sigma^F, \sigma^V\}$. For the I-O systems integrated under a particular set of tuning parameters, we calculate a statistic to evaluate the performance of the estimates with respect to each of the three economic properties. The I-O system with the statistics closest to their theoretical values is considered optimal.

While the descriptions of the exercises below start by providing some intuitive explanations for the desirable economic properties, they do not derive the desired properties explicitly from first principles. Instead, the exercises are meant to exhibit that methodological choices made to construct economic data from underlying statistics may be grounded in economic theory. Further, the exercises are conducted using readily available data. In the descriptions some suggestions are made for improvements in the empirical application in future research.

8.4.1   Equalization of Returns to Capital

The idea that returns to capital should be equalized across industries is straightforward. Simple arbitrage requires industries with below-average returns to sell their capital to industries with above-average returns to take advantage of the more profitable activity. Of course, if capital cannot be changed instantaneously because of adjustment costs, a putty-clay technology, or the quasi-fixity of capital, then the simple arbitrage argument breaks down. The fact that we do not estimate equalized capital returns under any calculation suggests that something more than data mismeasurement is needed to explain cross-sectional variation in capital returns. Nonetheless, data mismeasurement probably widens the distribution of returns; estimates that minimize the variation are indicative of an optimal combination of the expenditure-side and income-side data with respect to this metric.

To measure the performance of each integrated estimate, we calculate the return to capital for each year. We exclude Government Enterprises, Miscellaneous Industries, and Not Domestic Production from consideration because there is no presumption of profit-maximizing behavior in these industries. For each year we calculate the variance of returns across industries and then average the variance over the 1977–2001 period.

The return to capital is defined as capital income divided by an estimate of the wealth stock. Capital income equals value added less compensation paid to all types of labor less noncapital taxes on production and imports plus government subsidies. These data come from Bartelsman and Beaulieu (2004) as adapted from the GDP-by-industry data. Compensation is adjusted to include an imputation for the labor income of the self-employed; as measured in the NIPAs this income is counted in proprietors' income.[7] Noncapital production taxes are composed mostly of sales taxes.

---

7. The BLS makes the same adjustment in its productivity and cost estimates.

Simply plugging in the data on compensation, taxes, and subsidies assumes that these components of income paid are not mismeasured. The compensation data, at least to employees, is probably better measured than profits, interest, and proprietors' income; nonetheless, the idea that all of the mismeasurement of income resides in capital income is simply a maintained hypothesis that is not pursued further.

Estimates of the wealth stock are calculated based on detailed BEA estimates of investment by industry and by asset type. Wealth stocks were calculated using the appropriate formula (Hulten 1990) that is consistent with the age-efficiency schedule used in Bartelsman and Beaulieu (2004). The BEA investment data are adjusted for each input-output estimate of total investment to the extent that estimated deliveries to private fixed investment differs from the original estimate in the NIPAs on which the detailed BEA data are based.

In future work this exercise could be improved upon in three directions. The most difficult improvement would to provide some indicator of differences across industries in intangible investments that are not captured in the official data. Clearly, high implied rates of return—for example, in the movies and recreational services sector—could reflect returns to intangibles. A proxy for the cross-industry variation in intangible investment could be constructed based on expenditures on R&D and advertising. Next, the risk premium required may vary across industries, reflecting both differences in the time-series variation of returns, and differences in creditworthiness of firms in the industry. A proxy for this could be computed by looking at average bond yields and equity returns of publicly traded firms by industry, and by assessing the share of industry output that is provided by publicly traded firms, other corporations, and the noncorporate sector.

Finally, a third method could improve the rate of return in our exercise, now computed for each industry by dividing capital income with the wealth stock. The implied rate of return that equates the sum across asset types of the user cost times the stock of each asset to the total capital income differs from our crude measure because it takes into account differences across industries and asset types in depreciation rates, expected capital gains, and taxes.

8.4.2   Orthogonality of Innovations to Total Factor Productivity

The idea that variation in GDP is driven by productivity shocks that are common across industries is a central tenant of real business-cycle theories. Opponents to this theory have generally held that the size of the aggregate shock required to generate business cycle variation is implausibly large; candidate sources for such aggregate shocks, such as the weather, appear to amount to little. Simply adding up idiosyncratic shocks leads to an aggregate productivity shock that does not equal exactly zero, but because

of the law of large numbers the aggregate is too small unless the sector-specific shocks are large.[8]

Inherent in the counterargument to real business cycle models is that industry total factor productivity (TFP) growth rates should be uncorrelated. With measurement error, however, TFP growth rates can be correlated, even if they are orthogonal in reality. The measurement error can be correlated if it involves an allocation error of a fixed aggregate across industries. If the measurement error affects industries differently and this is somehow related to the business cycle—perhaps due to whether the product is a good or service—mismeasurement can also generate a correlation.

Economists have tested whether there is a common factor to industry productivity shocks (Lebow 1990; Forini and Reichlin 1998). In this exercise we do the opposite: we assume that this common factor is small and look for what combination of data produces a set of TFP growth rates that are as close to orthogonal as possible. To measure the orthogonality of TFP growth rates, we model the TFP growth rates as a linear function of a reduced number of principal components. The sum of the largest handful of standardized eigenvalues is a measure of the percent of the variation explained by the corresponding principal components; the smaller this measure, the more uncorrelated the TFP growth rates are.[9]

Industry TFP measures are calculated by modeling real gross output as a function of capital services, labor hours, and real intermediate inputs, using the usual Divisia formulation. Deflators for gross output come from the BEA's GDP-by-industry data set, as adopted in Bartelsman and Beaulieu (2004). The same gross output deflators are used to generate a deflator for intermediate input usage. Industry data on hours and capital services also come from Bartelsman and Beaulieu, although capital services built from investment flows are adjusted for differences in estimated aggregate deliveries to business fixed investment, as in subsection 8.4.1.

8.4.3   Stability of Intermediate Block

The idea that the coefficients of an I-O table should be stable is common in the literature. After all, the coefficients represent the structure and technology of an economy that evolve slowly due to "technical progress, exhaustion of natural resources, or variation in consumers' tastes"; the stability of the structure of the economy stands in contrast to final demand,

8. Horvath (2000) shows that the law of large numbers has to be augmented by the I-O structure of the economy. If the I-O table is sparse, then the law of large numbers applies at a much slower rate than is commonly presumed.

9. The fact that we compare twenty-one series with twenty-four years of data makes the measurement of orthogonality difficult. If the number of years in our data set was large relative to the number of series, we could choose a simpler measure, such as the determinant of the cross-correlation matrix.

which is less stable (Leontief 1953). Immediately, the question arises whether the stability of I-O coefficients should be measured using nominal data or real estimates (see Sawyer 1992 and references therein), and whether the values in the intermediate block should be constant with respect to the gross output of the supplying industries or the gross output of the demanding industries. De Mesnard (2002) uses the relative stability of the cells of the intermediate block divided by supplying industries versus those divided by demanding industries as a measure of whether an industry is "supply oriented" or "demand oriented."

For each estimate of the I-O system, we make four different calculations: two use nominal data, and two use real data, which are calculated by dividing the rows of the I-O table by the gross output deflators from Bartelsman and Beaulieu (2004). When using deflated measures, we ignore the obvious complications of taking ratios of chain-aggregated deflated data (Whelan 2000). Let $\mathbf{D(Y)}$ denote a square matrix with the gross output vector $\mathbf{Y}$ along the main diagonal and zeros otherwise. $\mathbf{I}$ is the intermediate block. Define allocation and technical coefficients thus:

$$\text{Allocation coefficients: } \mathbf{A}_t = \mathbf{D(Y}_t)^{-1} \cdot \mathbf{I}_t$$

$$\text{Technical coefficients: } \mathbf{T}_t = \mathbf{I}_t \cdot \mathbf{D(Y}_t)^{-1}.$$

We then take the standard deviation of each cell of $\mathbf{A}_t$ and $\mathbf{T}_t$ across time and then collapse this matrix into a single statistic by taking a weighted average of the standard deviations of each cell, where the weights equal the average of the absolute value of the cells of $\mathbf{I}$ over time.

8.4.4   Results

Figure 8.7 plots the results of these exercises. On the bottom axis of each panel are the values of $\{\sigma^F, \sigma^V\}$, displayed as σ^F on top of σ^V. Two other integrated I-O systems were calculated, denoted as $\{0, \infty\}$ and $\{\infty, 0\}$. The first system, $\{0, \infty\}$, is calculated by sweeping the vector $\mathbf{Y}\text{-}\mathbf{F}$ across the columns of the initial estimates of $\mathbf{I}$ without any reference to the initial values of $\mathbf{V}$; the value of $\mathbf{V}$ is calculated as a residual according to the value added identity. The second system, $\{\infty, 0\}$, is calculated by sweeping the vector $\mathbf{Y'}\text{-}\mathbf{V}$ across the rows of the initial estimates of $\mathbf{I}$, ignoring the initial values of $\mathbf{F}$; the resulting value of $\mathbf{F}$ is calculated using the gross output identity.[10]

The upper left panel plots the average cross-sectional standard deviation of the return to capital. Except for the estimate $\{0, \infty\}$, this measure of variation in the return to capital lies in the range 37.3 to 43.3. The minimum at 37.38 is at $\{.7, .3\}$, but 37.44 at $\{0, .5\}$ is also fairly close to the minimum. None of the data sets controlled to the GDI data ($\sigma^V = 0$) perform relatively well on this score.

---

10. Using the notation above, where $\mathbf{j}$ is a vector of ones that conforms to ($\bar{\mathbf{I}}$): $\mathbf{I}\{0, \infty\} = \mathbf{D}(\bar{\mathbf{Y}} - \mathbf{F}) \cdot \mathbf{D}(\bar{\mathbf{I}}\mathbf{j})^{-1} \cdot \bar{\mathbf{I}}$ and $\mathbf{I}\{\infty, 0\} = \bar{\mathbf{I}} \cdot \mathbf{D}(\bar{\mathbf{I}}'\mathbf{j})^{-1} \cdot \mathbf{D}(\bar{\mathbf{Y}} - \bar{\mathbf{Y}}')$.
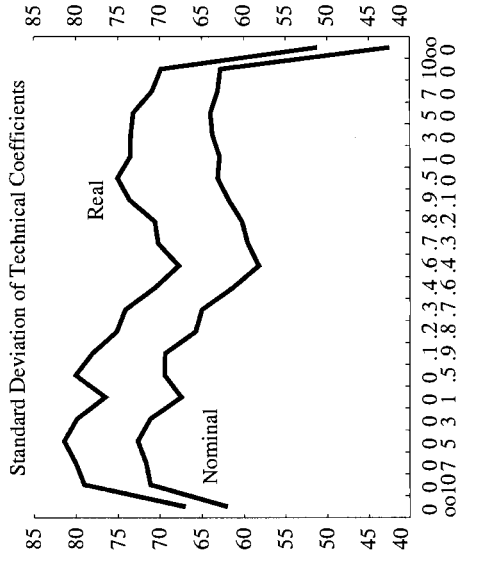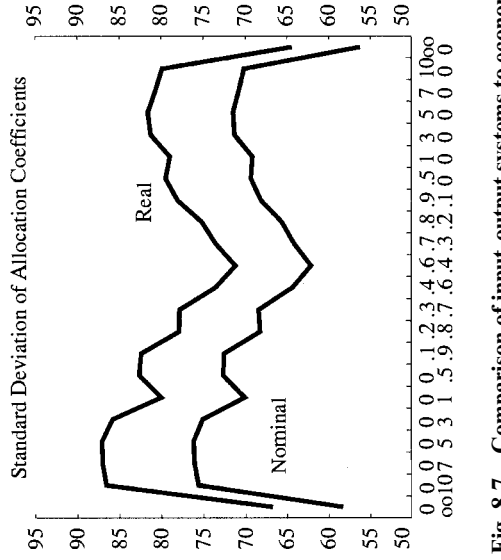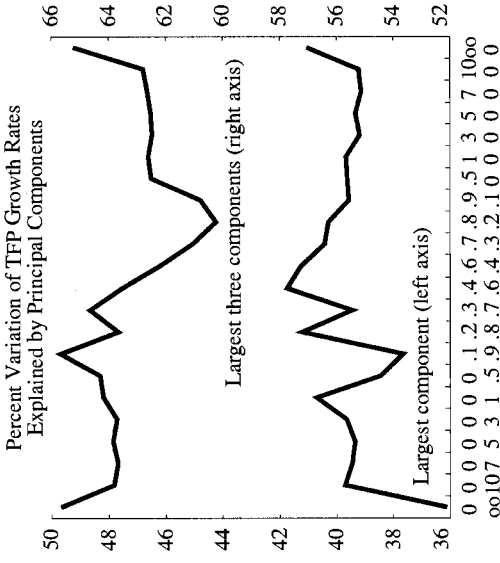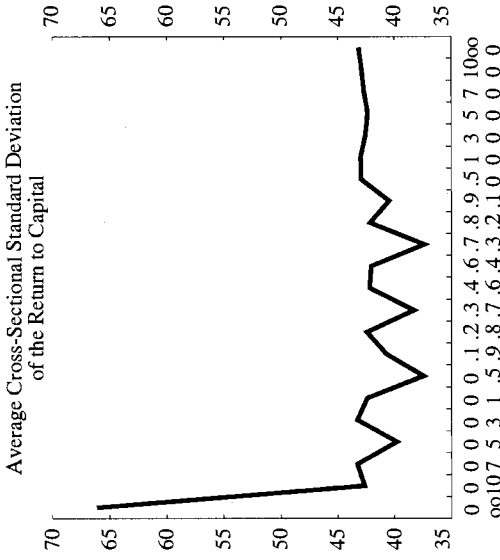
Fig. 8.7 Comparison of input-output systems to economic theory

Notes: Bottom axis shows tuning parameter for final demand (top row) and for value added (bottom row).

The upper right panel plots the percent of the variation of TFP growth rates explained by the largest principal component and by the largest three components. One principal component explains somewhere between 36 and 41 percent of the cross-sectional variation of TFP growth rates, with $\{0, \infty\}$ proving to be the least explainable among the integrated data sets calculated. However, using only one component to measure orthogonality is probably too restrictive, and we also present results using the three largest principal components. On this measure $\{0, \infty\}$ performs the worst, while $\{.8, .2\}$ at 60.25 percent has the least amount of variation explained by three principal components. Raising the number of components to four or five does not change this result, while adding even more components yields statistics that vary little across data sets.

The bottom two panels present results on the stability of the I-O coefficients. The bottom left panel plots the standard deviations of the real and nominal allocation coefficients; the bottom right panel plots the same for the technical coefficients. Excluding the tails, $\{0, \infty\}$ and $\{\infty, 0\}$, the data set with the most stable coefficients is $\{.6, .4\}$, with other data sets that roughly evenly mix the expenditure and income-side data also performing relatively well. The fact that the $\{0, \infty\}$ and $\{\infty, 0\}$ estimates produce the least variation in the standard deviation of real technical coefficients is essentially by construction because the calculation of the initial values of **I** are developed under the assumption that the real technical coefficients are constant. The stability of the nominal technical coefficients and the nominal and real allocation coefficients also benefit by this construction.

Taking the results together, the differences across data sets are not large, and some of the results do not smoothly vary when the data sets are ordered by tuning parameters. Nonetheless, they appear to point in a consistent direction: data sets constructed by mixing the information from the expenditure side and income side without controlling the aggregate to equal GDP or GDI yields estimates that perform well on all three criteria. The results also seem to favor a small bias towards the income-side data, a result that echoes Weale (1992).

8.5   Conclusion

In this chapter we employ industry estimates of deliveries to final demand and value added to investigate possible sources of the statistical discrepancy. We find that the expenditure-side data and the income-side data imply two different paths for the production of goods and services from the Machinery and Instruments, Trade, and Finance and Insurance industries that appear to be related to the statistical discrepancy. Important for the measurement of recent movements in productivity, there is an anomalous shortfall in 2001 in the change in private fixed investment implied from the income-side data relative to that measured from the expenditure-side data,

due to mismeasurement in sectors that include the high-tech industries. At a minimum, it might be useful to push on the source data for these industries to see if some improvement in data collection could help reconcile these discrepancies.

Our analysis also uncovered some other possible discrepancies that warrant some attention, even if they are not consistently related to the aggregate discrepancy. There are some important differences in our two sets of estimates of deliveries to final demand in the Chemicals, Refining, and Rubber and Plastics industry and in the Communications industry. There are also some significant differences in the estimates of value added in the Mining and Health Services industries.

Viewed differently, most of the statistical discrepancy shows up in PCE, but problems in the Machinery and Instruments industry also affect the statistical discrepancy in private fixed investment.

As a necessary step of this analysis we produced a consistent, integrated set of estimates of industry gross output, deliveries to final demand, intermediates used, and value added. We also produced a series of estimates and offered some means to judge how they should be combined. Some combination of the expenditure-side and income-side data should be employed, perhaps weighted more to the GDI data than to the GDP data.

We could not have written this paper if the BEA had not produced the wealth and the variety of the data that it does. Besides all of the information provided in the NIPAs, the GDP-by-industry data, and the published I-O tables, the importance of various estimates that the BEA makes available on its web site for researchers, such as the tables on underlying expenditure detail and the estimates of investment by industry and by asset type should not be overlooked. Of course, there would be no point in writing this paper if the BEA did not publish two estimates of domestic product; some countries only produce one estimate by balancing the information from expenditure-side and income-side data. If the BEA published only one estimate of domestic product, then only the BEA could have done the forensic analysis in this paper.

Even though "the man with two watches is never quite sure what time it is," the man with one watch may not realize that his watch has slowed or even stopped. An English version of this proverb that we have seen starts with "It's possible to own too much . . ."; as economists we know this cannot be true—especially with respect to data. Policymakers found important clues in the income-side measures of the transition of the economy when the production of and investment in high-tech goods pushed the growth rate of potential GDP higher (Jorgenson and Stiroh 2000).

As part of its strategic plan, the BEA has now published integrated value added I-O accounts with GDP-by-industry accounts. These integrated data add to GDP (Lawson et al., chap. 6 in this volume); they supplant the former GDP-by-industry data that add to GDI. While a published, con-

sistent, integrated data set that relates gross output, value added, and deliveries to final demand by industry is certainly useful, it comes at a cost. The new GDP-by-industry data are inconsistent with the prior data because the data now aggregate to GDP instead of GDI.

It is easy to recommend that others find resources in their budgets to provide additional data. Fortunately, the BEA already publishes a lot of the data that would be needed to develop a set of industry estimates of value added that add to GDP. In section 6 of the NIPAs—Income and Employment by Industry—the BEA provides data on the various components of income paid by industry. As discussed earlier, the problem with using these data directly is that some of the data are organized on a firm basis, instead of an establishment basis. However, if the BEA were to make available on its web site the factors that it uses to convert the data on a firm basis to an establishment basis—something the BEA will have to develop in-house anyway in order to prepare its integrated accounts—the research community could develop a second, consistent data set in real time that could be used to monitor and investigate future data discrepancies.

Appendix

Table 8A.1

Industry	SIC 87	Description
Agriculture	01–09	Farms, agricultural services, forestry, fishing, hunting, and trapping
Mining	10–14	Metal mining, coal mining, oil and gas extraction, and mineral mining
Construction	15–17	Construction
Wood, furniture, paper, and printing	24–27	Manufacturers of lumber and wood, furniture, paper, and printing
Primary durable manufacturing	32–34	Stone, clay and glass, primary metal, and fabricated metal manufacturing
Machinery and instruments	35–36, 38–39	Machinery, electrical machinery, instruments, and miscellaneous manufacturing. This industry includes computers, communications equipment, and semiconductors.
Transportation equipment	37	Motor vehicles and parts, aircraft and parts, and other transportation equipment
Food and tobacco	20–21	Food and beverages and tobacco manufacturing
Textiles, apparel, and leather	22–23, 31	Textiles, apparel, and leather manufacturing
Chemicals, refining, and rubber and plastics	28–30	Chemicals, petroleum refining, and rubber and plastics manufacturing
Transportation	40–42, 44–47	Trucking, water, rail, and air transport, warehousing, pipelines (except natural gas), and transportation services
Communications	48	Telephone and telegraph, radio and television, and other communications services
Utilities	49pt.	Electrical, natural gas, and water and sanitary services utilities. It excludes government enterprises such as TVA and Bonneville.
Trade	50–59	Wholesale and retail trade
Finance and insurance	60–64, 67	Depository and nondepository institutions, securities dealers and brokers, insurance carriers and agents, and holding companies
Real estate	65	Real estate, excluding imputations for owner-occupied housing and the rental value of nonprofits' capital[a]
Hotels and other lodging	70	Hotels and other lodging
Personal services	72, 75–76	Personal services, automotive repair services and parking, and miscellaneous repair services
Business services	73	Business services, including software and data processing
Movies and recreation services	78–79	Motion pictures, and amusement and recreation services
Health services	80	Health services
Legal services	81	Legal services
Other services	82–84, 86–87, 89	Social services, museums, membership organizations, engineering, accounting, research and management services, and miscellaneous services
Government enterprises	43, 49pt, other	Federal, state, and local government enterprises, including the postal service, TVA, and Bonneville Power
Miscellaneous industries	88, other	Private households, owner-occupied housing, and general government
Not domestic production	—	Imports, used and secondhand goods, and scrap

[a]The rental value of nonprofits' capital equipment and structures was distributed to other industries according to estimates of nonprofit activity in those industries.

References

Bacharach, Michael. 1970. *Biproportional matrices and input-output change.* Cambridge, UK: Cambridge University Press.

Baker, Dean. 1998. The new economy does not lurk in the statistical discrepancy. Challenge 41 (4): 5–13.

Bartelsman, Eric J., and J. Joseph Beaulieu. 2003. Techniques to reconcile data with linear constraints. Washington, DC: Board of Governors of the Federal Reserve System. Mimeograph, January.

———. 2004. A consistent accounting of U.S. productivity growth. FEDS Working Paper no. 2004-55. Washington, DC: Board of Governors of the Federal Reserve System, October.

Bishop, Y. M., S. E. Fienberg, and P. W. Holland. 1975. *Discrete multivariate analysis: Theory and practice.* Cambridge, MA: MIT Press.

Bosworth, Barry P., and Jack E. Triplett. 2003. Services productivity in the United States: Griliches' services volume revisited. Washington, DC: The Brookings Institution. Mimeograph, September. Available at http://www.brookings.org/views/papers/bosworth/20030919.htm.

Chen, Baoline, and Dennis J. Fixler. 2003. Measuring the services of property-casualty insurance in the NIPAs: Changes in concepts and methods. *Survey of Current Business* 83 (10): 10–25.

Congressional Budget Office (CBO). 2003. CBO's economic forecasting record: An evaluation of the economic forecasts CBO made from January 1976 through January 2001. Available at http://www.cbo.gov/showdoc.cfm?index=4639&sequence=0.

Council of Economic Advisers. 1997. Economic report of the president. Washington, DC: Government Printing Office.

De Mesnard, Louis. 2002. Forecast output coincidence and biproportion: Two criteria to determine the orientation of an economy; Comparison for France (1980–1997). *Applied Economics* 34:2085–91.

Dean, Edwin R., Michael J. Harper, and Phyllis F. Otto. 1995. Improvements to the quarterly productivity measures. *Monthly Labor Review* 118 (10): 27–32.

Edelstone, Marc, Louis Gerhardy, Harlan Sur, Steven Pelayo, and Jay Iyer. 2003. Transition to 300-mm wafers should drive secular changes. Equity research. New York: Morgan Stanley, December.

Fixler, Dennis J., Marshall B. Reinsdorf, and George M. Smith. 2003. Measuring the services of commercial banks in the NIPAs: Changes in concepts and methods. *Survey of Current Business* 83 (9): 33–44.

Forini, Mario, and Lucrezia Reichlin. 1998. Let's get real: A factor analytic approach to disaggregated business cycle dynamics. *Review of Economic Studies* 65 (3): 453–73.

Greenspan, Alan. 2004. Risk and uncertainty in monetary policy. Remarks at the meeting of the American Economic Association. 3 January, San Diego, California.

Günlük-Şenesen, G., and J. M. Bates. 1988. Some experiments with methods of adjusting unbalanced data matrices. *Journal of the Royal Statistical Society* Series A (*Statistics in Society*) 151 (3): 473–90.

Horvath, Michael. 2000. Sectoral shocks and aggregate fluctuations. *Journal of Monetary Economics* 45:69–106.

Hulten, Charles R. 1990. The measurement of capital. In *Fifty years of economic measurement studies in income and wealth,* ed. Ernst R. Berndt and Jack E. Triplett, 119–52. Chicago: University of Chicago Press.

Jorgenson, Dale W., and Kevin J. Stiroh. 2000. Raising the speed limit: U.S. eco-

nomic growth in the information age. *Brookings Papers on Economic Activity,* Issue no. 1:125–211.

Klein, L. R., and J. Makino. 2000. Economic interpretations of the statistical discrepancy. *Journal of Economic and Social Measurement* 26 (1): 11–29.

Lebow, David E. 1990. The covariability of productivity shocks across industries. Board of Governors of the Federal Reserve System, Economic Activity Section, Working Paper no. 104. Washington, DC: Federal Reserve.

Leontief, Wassily. 1953. *Studies in the structure of the American economy: Theoretical and empirical explorations in input-output analysis.* New York: Oxford University Press.

Moulton, Brent R. 2000. Getting the 21st century right: What's underway? *American Economic Review: Papers and Proceedings* 90 (2): 253–58.

Nordhaus, William D. 2000. New data and output concepts for understanding productivity trends. Cowles Foundation Discussion Paper no. 1286. New Haven, CT: Yale University, November. Available at http://cowles.econ.yale.edu/P/ab/a12/a1286.htm

Parker, Robert P., and Eugene P. Seskin. 1997. The statistical discrepancy. *Survey of Current Business* 77 (8): 19.

Sawyer, John A. 1992. Forecasting with input-output matrices: Are the coefficients stationary? *Economic Systems Research* 4 (4): 325–48.

Schneider, M. H., and S. A. Zenios. 1990. A comparative study of algorithms for matrix balancing. *Operations Research* 38:439–55.

Stone, Richard, D. G. Champernowne, and J. E. Meade. 1942. The precision of national income estimates. *Review of Economic Studies* 9 (2): 111–25.

U.S. Bureau of Economic Analysis. 2001. A guide to the NIPAs. Washington, DC: Bureau of Economic Analysis, Department of Commerce. Available at http://www.bea.doc.gov/bea/an/nipaguid.htm

Weale, Martin. 1992. Estimation of data measured with error and subject to linear restrictions. *Journal of Applied Econometrics* 7 (2): 167–74.

Whelan, Karl. 2000. A guide to the use of chain aggregated NIPA data. FEDS Working Paper no. 2000-35. Washington, DC: Board of Governors of the Federal Reserve System, June.

Yuskavage, Robert E., and Erich H. Strassner. 2003. Gross domestic product by industry for 2002. *Survey of Current Business* 83 (5): 7–14.