Chapter 1

Introduction

When a country experiences a balance of payments problem, the typical remedy mix proposed by the International Monetary Fund consists of fiscal austerity, tight monetary policies, devaluation, privatization, elimination of subsidies and trade liberalization, combined with low interest rate loans.

Throughout the late 1980s, Ghana has been hailed as a success story for that policy mix, known internationally as "neoliberal policies" and in Ghana as the "Economic Recovery Program" or ERP. However, Ghana's performance has been increasingly disappointing during the 1990s, with economic growth slowing down and government deficit surging out of control.

This thesis explores the reasons for that slowdown, its distributional implications, and the extent to which the behavior of the Ghanaian economy conformed to commonly used assumptions in economic models of developing countries.

The following section states the specific research questions of this thesis. Having stated the questions, we proceed to discuss the methods we use in answering these research questions, as well as the intermediate, technical results of these methods, in Section 1.2. Section 1.3 then uses these intermediate results to answer our research questions, and Section 1.4 discusses some limitations of our approach.

1.1 Research Questions

In this dissertation, we address three closely interlinked research questions.

1.1.1 Why Was Macroeconomic Performance Not Sustained?

In the 1990s, what caused

1. Balance of Payments problems
2. Inflation
3. Low GDP growth

1.1.2 What Were the Distributive Impacts?

How did the structural adjustment policies affect the real incomes of different population groups during the 1990s?

1.1.3 Does the behavior of the Ghanaian Economy in 1990s conform to the assumptions of economic models?

Once we understand the key causal mechanisms for the Ghanaian economy in the period under study, we are in a position to evaluate the common assumptions of models and qualitative descriptions that are used to assess impacts of structural adjustment programs.

1.2 Methods and Intermediate Results

Our research for this thesis consists of three steps. First, we compile a stock-flow consistent time series of yearly Social Accounting Matrices and monthly Financial Accounting Matrices for 1990-2001. Then, we use this dataset along with additional monthly macro time series (such as interest rates) to statistically test a series of key relationships, which we have selected based on Computable General Equilibrium (CGE) and macroeconomic theory, so as to achieve a systematic picture of the Ghanaian economy's structure. Finally, the results of these investigations are used in qualitative discussion to address our research questions. The following sections describe each step in more detail.

1.2.1 Compile a Dataset: SAM/FAM time series for 1990-2001

We begin by compiling a yearly time series of Social Accounting Matrices (SAMs), each of which is a complete picture of the money flows in the Ghanaian economy in a given year, at a certain level of aggregation. Then, we augment this by a monthly time series of Financial Accounting Matrices (FAMs) that describe the financial stocks held by the various institutions in the Ghanaian economy (the SAM and FAM methodologies are discussed in detail in Chapter 5). The two datasets are logically connected as the net lending flows of an institution in each SAM have to be consistent with the changes in that institution's net worth (after correcting for revaluation). We achieve that consistency by forcing the SAMs to yield the let lending flows implied by the FAMs, for the following reason.

In developing countries, financial asset stock time series are available and generally more reliable than most flow data. Besides being interesting in its own right, knowledge of financial asset stock time series allows us to get more reliable estimates for net financial savings (net lending) of each institution. In an economy where some important sectors are demand-driven (that is, probably, any economy), knowledge of net lending is important to account for injections and leakages contributing to the aggregate demand. Unfortunately, net lending is normally computed as a residual from flow data, and is thus quite unreliable; using asset stocks for that purpose makes for a much more reliable estimate.

Compiling a SAM/FAM time series dataset such as ours is worthwhile for several reasons: firstly, in the process of compilation one discovers the errors and inconsistencies in the data and can to an extent correct for these, or at least draw conclusions as to which uses of that data are meaningful even in view of the errors and which are not (Chapter 7 discusses several grave errors in official data that we could discover and to an extent correct).

Thirdly, a sufficiently long stock-flow consistent SAM-FAM time series provides a sufficient data base for estimation *and* validation of a multisectoral dynamic model. A complete set of real and nominal flow time series allows us to derive the causal behavior of the model from history, instead

of specifying it a priori as the CGE models do, be they neoclassical or structuralist. Further, if we split the SAM-FAM time series into an estimation and a validation segment, we can also gain an idea of how good the chosen causal structure is at predicting future behavior of the economy. Such an approach unites the advantage of a CGE-like ability to handle sectoral disaggregation and a high number of variables with an ability to specify measures of confidence of the model output.

Finally, as the present project demonstrates, such a dataset can be used to understand the overall structure of an economy through statistical testing of specific relationships, even without estimating a complete model. This is an advantage for our approach's potential to be replicated in other developing countries, where technical capacity constraints are often binding.

Summing up, we demonstrate that in spite of low data quality and availability, it is feasible to construct stock-flow consistent SAM/FAM time series for Ghana for more than a decade, providing a strong foundation for further analysis. While that is not a novel approach for industrialized countries with their high data availability, to the best of our knowledge this is the first time such a dataset has been compiled for a developing country.

1.2.2 Statistical Analysis of Selected Variables

The statistical analysis we undertake consists of two parts. On the one hand, macro time series such as the Consumer Price Index (CPI), interest rates, etc. are available monthly, so that we can use reasonably sophisticated statistical techniques. On the other hand, the real side and sectoral data in only available yearly (for a total of 12 data points in our time period), necessitating a different strategy. Let us discuss each in turn.

We use ARIMA-X regressions to explain the behavior of the three variables Consumer Price Index, broad money supply, and the exchange rate, which we denote (in log form) `cpi`, `m2` and `er`. All of these are I(1), and in none of the initial regressions was cointegration observed. Therefore, the regressions are run on the first differences of the logs (corresponding to growth rates of the underlying variables).

We start out by investigating the behavior of the Consumer Price Index, arriving at both expected and unexpected results. On the expected side, money supply is highly significant for price formation, and takes its effect gradually over with a lag of two to nine months. Increases in fuel and wholesale food crop prices are also highly significant, but their impacts happen over a much shorter period (0-2 months). Using just these three variables, we could predict `cpi` surprisingly well four years into the future (1998-2001), with coefficients estimated using pre-1998 data only.

On the surprising side, exchange rate depreciation, interest rate changes, and GDP growth appear to not have a significant direct effect on inflation.

As we find that growth of broad money supply is the major driver of CPI inflation, we proceed to investigate the dependence of broad money supply growth on the changes in interest rates, inflation, exchange rate depreciation, and base money growth, as well as its own past values. The results show a strong dependence on the exchange rate, monetary base, and Consumer Price Index, as well as a significant but rather small response to interest rates.

From a .37 coefficient of the depreciation rate in the `m2` regression, combined with the fact that exchange rate depreciation was not significant in the `cpi` regression, we could draw the somewhat surprising conclusion that the main channel for the impact of currency depreciation on inflation is not through cost-push such as cost of imports, but through revaluation of the money supply. In contrast to the CPI regression, all impacts on money supply are quite fast, taking at most 3 months for the full impact to be felt; the predictions of post-January 1998 values based on the regressions

exhibit realistic behavior and are quite robust with respect to the estimation time period, but not as precise as the CPI predictors, tending to over-estimate broad money supply.

The two regressions we just discussed imply that if we consider broad money supply `m2` to be given in advance (as a function of time), a change of one percent in `m2` would over time lead to a change of about .33 percent in `cpi`; conversely, if we considered the price level to be exogenous, a one percent increase in it would lead over time to about a .67 percent increase in `m2`. In reality, neither of them is given, but rather both evolve (approximately) according to the behavioral equations that we have estimated in the previous section. Therefore, the two variables form a feedback loop - a change in the money supply, happening for whatever reason, will lead to an increase in the price level, which will in turn lead to an increase in money supply, etc., amplifying the initial impact. The total impact of the feedback loop can be computed as $1/(1 - 0.33 * 0.67) \approx 1.3$, so that the feedback loop amplifies by about 30% any effects of other variables on CPI inflation or broad money supply growth.

Following the same model identification strategy that led us to success in explaining the behavior of `cpi` and `m2`, we identify a model of exchange rate behavior. The only two significant variables turn out to be the import price index, each 1% increase in it translating into a 4% depreciation over a year's time; and the interest rate, with interest rate increase by 1% *increasing* depreciation by one third of a percent. Unlike in the cases of `cpi` and `m2`, however, the estimated model proves to be quite bad at predicting values of `er` when used as a recursive equation.

These problems, together with a brief qualitative discussion of exchange rate behavior, lead us to conjecture that better understanding of the exchange rate behavior would require separate study of managed floating and freely falling periods, as well as of the conditions for the change between the two modes.

While we are lucky enough to have monthly time series for the key nominal indicators (such as the Consumer Price Index, the exchange rate, and the money supply), the real-side sectoral data, along with most nominal flow data, is only available on a yearly basis.

As our sample length of 12 years is too short to allow for rigorous econometrics, we use a different method to evaluate hypotheses of real-side behavior. We take the years 1990-1997 as the estimation period and the years 1998-2001 as the validation period. We use the estimation period data to estimate the coefficients of behavioral hypotheses and then use their ability to predict the validation period data as the way to choose between the different hypotheses.

This approach has the potential to increase our confidence in the choice of functional form, as a good fit in the estimation period does not necessarily translate into a good prediction ability in the validation period. On the other hand, if a functional formulation was good at predicting the validation period data when estimated with only estimation period data, we can have some confidence in its predictions for the future (or counterfactuals in the past) when estimated using the data of both periods. After having thus identified the functional form, we use ordinary least squares to produce standard errors for the intercepts and elasticities involved.

The first relationship we test is the venerable constancy of the Incremental Capital Output Ratio, that is the hypothesis that output growth is directly proportional to investment. In line with previous research, this hypothesis is not supported by the data.

We then turn to the productive sectors. In agriculture, we observe a price elasticity of substitution between export and nontraded crops of .42(.15); until 1995, land surface under cultivation is approximately constant, and the wholesale price of food crops relative to CPI moves in the same direction as food crop production, suggesting a demand-driven regime. However, starting in 1995, land surface under cultivation begins to grow substantially, and the relative price of wholesale

food crops moves consistently in the opposite direction to food crop output, suggesting a shift to a supply-constrained mode.

In industry, the elasticity of substitution between exports and nontraded output turns out to be negative, equaling -1.00(.23). This is not compatible with a neoclassical productivity-frontier description, but fully compatible with a demand-driven nontraded industrial sector. The hypothesis that nontraded industrial output is demand driven is further supported by the behavior of said elasticity when we change the estimation time period or exclude mining from the estimated ratio, as discussed in Chapter 9. Thus, we conclude that the nontraded industrial sector is demand-driven.

Turning to the demand side, we discover that the government was the major source of demand injections throughout the period under consideration. As nontraded industry is demand-driven, and most likely so is the informal sector, we conclude that government demand was one of the key drivers of the economy.

As far as investment is concerned, we see that investment as a share of GDP has been low (about 15%) and declining; investment demand decreased by about 2% (not percentage points) for each percentage point increase in interest rates.

Finally, turning to import demand, we find a quite large price elasticity of substitution between nontraded and imported manufactures, namely -1.18(.31). This implies that the balance of payments will be quite sensitive to both exchange rate and tariff changes. The corresponding elasticity for services was -0.38(0.11), smaller but still substantial.

1.2.3 Discussion of Quantitative Results' Implications

Having thus reached a number of quantified insights into the structure and behavior of the Ghanaian economy, we proceed to use these to answer our research questions. We do so by means of a verbal discussion, rather than an overall formal model of the Ghanaian economy. The latter decision is defended in Section 1.4.1; let us now proceed to the implications of our data investigations for our research questions.

1.3 Results

1.3.1 Neoclassical/Monetarist vs. Structuralist Theory

The estimates of key macro and sectoral relationships we have discussed above, such as the various elasticities and the determinants of inflation and money supply growth (all of them complete with significance measures), allow a data-driven comparison of the applicability of the two competing schools, neoclassical/monetarist and structuralist, to the Ghanaian economy.

On the macro/nominal side of things, we observe a very strong link from money supply growth to inflation (a monetarist staple assumption) and a weaker but also significant link from inflation to money supply growth (a common structuralist theme). The two variables thus form a feedback loop, which amplifies by about a third all effects of other variables on either inflation or money supply. While broad money supply growth is the dominant determinant of inflation, the price of fuel and the wholesale price of food crops are also important. Additional influences are growth of monetary base and exchange rate depreciation, both of which impact inflation not directly, but through their effects on broad money growth.

Overall, on the nominal side the monetarist hypotheses are better supported by the data than their structuralist counterparts, but with two important caveats. Firstly, certain cost factors

were also important, though less so than money supply growth; and secondly, the relevant money supply measure was broad money, and not monetary base that is commonly used for that purpose in monetarist CGE models.

On the real side, agriculture is supply constrained while industry is demand-driven - a setup quite common in structuralist models, as opposed to the neoclassical assumption that all sectors are always at full capacity utilization and therefore supply-constrained. On the other hand, import demand is quite sensitive to the relative price of imports vs. nontraded goods, which is a typical feature of neoclassical models, while structuralist models often assume that quantity effects dominate and relative price effects are of moderate importance.

Summing up, on the side of price formation monetarist theory has the upper hand, while on the real/sectoral side of things structuralist theory is better supported by the data. In both cases, however, the optimal description must include elements of both theories.

1.3.2 Reasons for Poor Macroeconomic Performance

The quantitative understanding of the Ghanaian economy that we have gained also allows us to give a well-founded explanation to the poor macroeconomic performance of the Ghanaian economy in the 1990s. Briefly, the story is as follows.

First of all, import liberalization strangulated domestic manufacturing, as the demand for manufactures was redirected towards imports. At the same time, withdrawal of government support programs for agriculture, combined with population growth, led to an increasing supply problem in nontraded food crops, which led to an increase of food crop imports, especially rice. As the growth in exports was not sufficient to pay for these increases in imports, a persistent balance of payments deficit appeared. As that deficit accumulated in a growing foreign debt, the resulting interest payments further worsened the balance of payments.

The next part of the story are the occasional bursts of government spending, particularly in the years around the first general election of the 1990s, in 1992-1994. While the demand generated thereby did stimulate domestic manufacturing somewhat, it also led to abnormal money supply growth and resultant inflation. Furthermore, a large part of that demand spilled into imports, resulting in even more balance of payments deficits. The inflation combined with balance of payments deficits led to exchange rate depreciation, which in turn boosted money supply growth and inflation in yet another positive feedback loop.

The demand strangulation in industry and the withdrawal of government support programs in agriculture are a part of the explanation for the low GDP growth rate in the 1990s as compared to the 1980s; another reason is that the starting point for GDP in early 1980s was so low that high initial growth rates were easy to realize - and this was no longer the case in 1990s. The final reason for low GDP growth are the high, and volatile, inflation rates. As the result of the the Bank of Ghana's attempts to keep real interest rates positive, nominal interest rates became quite high, and real interest rates quite volatile. That further depressed investment demand, adding to the effects of low demand.

1.3.3 Distributional Impacts

As this thesis focuses on analyzing macro and sectoral relationships, it can unfortunately only shed a limited amount of understanding onto the distributional implications of the government policies and external shocks during the 1990s. What we can do is to provide an understanding of how

incomes of the various sectors (at our level of aggregation) were affected, which is a necessary first step for a more detailed analysis.

The overall impacts of the government policies certainly benefited export producers, which in Ghana means mainly mining and cocoa. As mining is to a large extent foreign-owned, those profits were not really passed on to the employees (resulting in a profit rate of 50% for mining companies in 1993, far above any other sector), so that the main beneficiaries were the cocoa farmers. These certainly did benefit, with the average cocoa farmer household's income exceeding the average food farmer household's income by about 50%. However, their income was still below that of an average formal sector (industry or government) employee household by another 50%; and geographical reasons limited the expansion of cocoa production.

All agriculture was badly hit by withdrawal of government support programs, and manufacturing was adversely affected by the leakage of demand into imports - thus the net incentives, apart from the benefits to cocoa producers, were towards relocation of labor into the informal sector, comprising a variety of transport, retail, and personal service activities, to a large extent untaxed.

1.4 Limitations

While we believe that this thesis provides significant insight into its research questions, it also was more successful in some areas than in others. This section's purpose is to mark out the main limitations or our work.

1.4.1 No Formal Overall Model

An important limitation of this thesis is that while we compile a dataset sufficient to estimate a complete model of the Ghanaian economy at a certain level of disaggregation (6 sectors, 5 institutions), we do not, in fact, estimate such a model.

In addition to compiling the dataset, we do estimate key relationships that would be useful for building such a model in Chapters 9 and 11, and present accounting frameworks that would form the backbone of such a model in Chapters 8 and 10. However, we stop short of putting all the pieces together and estimating a comprehensive model of the Ghanaian economy in the 1990s. How do we justify this?

There are four major reasons we chose that course of action, namely transparency, management of uncertainty, optimal return to effort, and replicability. Let us discuss each in turn. By management of uncertainty we mean that as the data for different parts of the system is of varying quality, the level of confidence of results derived from that data will also vary. While with monthly data we can test a number of hypotheses involving a variable, yearly data at best allows us to get a reasonable estimate of something like one elasticity per time series; and one or two time series had to be derived as pro-rating from point estimates, making them even less reliable. In a model, all these would appear as numbers with little to indicate their relative reliability - while in a qualitative discussion such as ours, it is possible to distinguish between them.

A similar reasoning applies to higher transparency: in a verbal discussion, it is much more apparent to the reader where the conclusions come from - the reasoning must be stated argument for argument; on the other hand, while model equations are a language in themselves (and arguably a more powerful and versatile one for certain purposes), it is a language that is spoken by a much smaller audience. Therefore, when addressing a non-technical audience, and presenting the results

as coming from a model, one risks to be perceived as an oracle for a black box - not an ideal communication strategy.

Neither of these problems is in itself fatal. One could cleanly combine data of varying quality in the same model by specifying model-wide confidence measures that take these quality differences into account. A model-wide optimization, followed by model-wide Monte Carlo experiments, could generate simultaneous confidence intervals for all parameters as well as model output. On the transparency side, nothing prevents a modeler, after having built a model, from also deriving the results in a verbal discussion; the results will stand all the stronger if they come from both a discussion and a modeling exercise.

That brings us to the last two reasons, namely return to effort and replicability. While model-wide estimation and optimization, followed by model-wide Monte Carlo experiments, would in fact allow us to effectively manage uncertainty, and additional discussion would clarify the results to a non-technical audience, the overall effort required would be quite high. That, in turn, would restrict the usefulness of the approach to the institutions that have the appropriate technical capacity and willing to make the extra effort, a fairly high threshold if one is talking of developing countries. Finally, even when the ultimate objective is building a comprehensive model, it is useful to be able to generate analyses like ours beforehand, both to map out the shape of the model-to-be and to generate intermediate output for funding purposes.

Summing up, we don't deny that estimating an overall formal model of Ghana in the 1990s would be very interesting and yield additional insights compared to our approach. However, our approach already surpasses the existing ones in being more systematic than isolated econometric regressions, and in providing measures of confidence in its results based on time series data, which CGE models can't do. It lays the necessary data foundation for building a formal model, and reaches substantial analytical and policy conclusions from study of carefully selected variables; thus, it is well suited for replication in other developing countries, where both varying data quality and limited technical capacity are common.

1.4.2 Limited Analysis of Distributional Impacts

Another important limitation of the present thesis is that it does not provide a deep analysis of distributional impacts of the government policies discussed. The reason for this is twofold, namely the scope of our effort and poor data availability.

The present thesis contains an analysis of the Ghanaian economy in the 1990s on the macro and sectoral levels, based on yearly or monthly time series. The macro/sectoral scope of the effort limited the amount of understanding we could provide of the finer distributional impacts of government policies. A good understanding of the sectoral structure of the economy is a necessary foundation for a sound distributional analysis, as it enables us to see how the incomes of broad occupation groups are affected by macro and sectoral policies and shocks. This is the kind of very broad analysis that we were able to conduct here.

Substantial advances in distributional analysis would require either a further increase in sectoral resolution of the data, or better the use of micro-level data such as household surveys. Finer sectoral price and output data would have been even harder to collect than our present dataset. On the micro data side, there are only two household surveys available for the period under study, and these are in form of raw tables with sparse documentation, requiring much effort to extract useful summary data.

Given the high effort of collecting and processing the necessary data for a deeper distributional

analysis, along with an already high workload from the other parts of the research project, we regretfully restrict ourselves to a quite shallow sectoral-level distributional analysis.

1.5 Structure of This Manuscript

The structure of this thesis is designed to address the diverse levels of technical background that its readers will have. The challenge is to make the document interesting to readers with a high level of expertise in economics, while at the same time keeping it accessible to readers with limited economics background.

We aim to achieve it by separating the document into "narrative" and "technical" chapters. Thus, a non-technical reader might want to read Chapters 2 to 6 that provide introduction to all relevant concepts and techniques, skim the technical details of Chapters 7, 8, 9, 10, and 11 and get all the key points from the discussion of these chapters' results in Chapters 12 and 13. Also, at one of the committee members' request, each chapter includes a short summary.

On the other hand, the readers with substantial economic expertise may wish to read Chapters 1 and 6 for a brief summary of the research question and the overall strategy of this thesis, skim or skip Chapters 2 to 5, and turn to Chapters 7, 9, and 11 for the empirical/data-related content and to Chapters 8 and 10 for some modeling technique issues.

Chapters 12 and 13 are meant for all readers as they pull together all the threads from the previous chapters, and Chapter 14 provides some concluding remarks.