

The Impact of US Subsidies on the World Cotton Market: A Reassessment

Ben Shepherd, Principal.

March 17, 2004.

THE IMPACT OF US SUBSIDIES ON THE WORLD COTTON MARKET: A REASSESSMENT

Ben Shepherd*

Groupe d'Economie Mondiale (GEM)
Institut d'Etudes Politiques de Paris
2, Square de Luynes
75007 Paris, France

This Version Dated: 17 March 2004

Abstract: Following a critical review of the existing quantitative literature on cotton subsidies, a vector autoregression (VAR) is used to model the effects of US subsidies on the world cotton market from 1965 to 2001. Surprisingly, subsidies are found to have only a limited impact on prices despite their effects on production and consumption. The dynamic relationships between quantities, prices, stocks and subsidies are found to be considerably more complex than those suggested by basic theory. Finally, simulation results indicate that even large reductions in US subsidies will not necessarily lead to significantly higher world prices.

JEL Codes: Q17; C32.

Keywords: Vector Autoregression; Subsidies; World Prices; Cotton Market

* The author is Research Assistant at the GEM and is currently completing his Ph.D. in economics jointly at Sciences-Po (Paris) and the Graduate Institute of International Studies (Geneva). Sincere thanks to Olivier Cattaneo (GEM), Prof. Patrick Messerlin (GEM) and particularly to Sébastien Miroudot (GEM) for helpful discussions and comments, as well as to Gérald Estur (ICAC), Richard Pazdalski (USDA) and Phil Sronce (USDA) for supplying data. Remaining errors are, of course, my own responsibility.

Contact: benjamin.shepherd@sciences-po.org.

“Through the subsidies it provides to American cotton farmers, the US Government is driving down world cotton prices and the household income of cotton farmers in Africa and elsewhere. The injury is ... real ... and it is being inflicted on some of the world’s poorest people by its wealthiest nation.” Oxfam (2002), p.6.

The issue of rich countries’ cotton subsidies erupted onto centre stage at the WTO in May 2003, when Benin, Burkina Faso, Mali and Chad (“the WCA countries”) presented their sectoral poverty reduction initiative (Benin et al., 2003) to the Special Session of the Committee on Agriculture. (For a general overview, see Goreux, 2004a).¹ They called for the abolition of those subsidies, which they said encouraged oversupply from relatively inefficient Northern producers, thereby lowering world prices and reducing the WCA countries’ export earnings. A similar allegation forms the basis of an on-going dispute between Brazil and the US over the latter’s subsidies to upland cotton producers (Brazil, 2002).

Given the political importance that cotton subsidies have acquired in the wake of the WCA countries’ submission, the somewhat limited attention they have received to date from empirical economists is surprising. While a number of studies attempt to model the possible impact on the world market of removing subsidies, there does not yet appear to be any serious analysis of the historical effects of subsidies, nor any attempt to extrapolate directly from that experience to model the impact of eliminating or reducing subsidies in the future.

This paper constitutes a preliminary attempt to fill that gap in the literature. In Section 1, I provide a critical review of past studies, focusing on those that the WCA countries relied on in framing their negotiating proposal. A number of important, but largely unstated, methodological limitations are highlighted. Section 2 outlines an alternative modelling strategy that is more strongly grounded in the analysis of historical experience. That approach

¹ My sincere thanks to Louis Goreux for supply me with copies of Goreux (2004a) and (2004b).

is implemented in Section 3 and some preliminary results are presented. Section 4 concludes and sketches out an agenda for further research.

1. The Empirical Support for the WCA Countries' Proposal

The economic logic underlying the sectoral initiative seems uncontroversial. The purpose of this section is to examine whether that logic is as well-grounded in the facts as it is in standard economic ideas. Before embarking on that analysis, it is useful to summarise the principal factual claims that were actually made (Benin et al., 2003):

- The WCA countries stated that they had lost US\$250m in export earnings for the crop year 2001-2002, a figure sourced from a previous version of Goreux (2004b).
- They suggested that in the absence of Northern subsidies, the world cotton price would have been 17 cents per pound higher in 2000-2001 and 31 cents per pound higher in 2001-2002, citing studies by ICAC that are reported in ICAC (2002).

The two claims are intimately related and both come back to the idea that in the absence of subsidies, world cotton prices would be significantly higher than they in fact are. The focus for the following literature review is therefore on the way in which previous work has sought to quantify that effect.

1.1. The Goreux (2004b) Model

In his report prepared at the request of the WCA countries' Ministers of Agriculture, Louis Goreux used a standard partial equilibrium model of the cotton market to quantify the impact of rich country subsidies. In summary form, his analysis proceeded as follows (see Figure 1 for a graphical representation):

- Farmers in non-subsidising countries receive the world price (p^w) for their cotton. Those in countries that subsidise receive the world price plus a per unit subsidy ($p^w + s$).

- Abolition of subsidies means that all farmers receive p^w , representing an initial fall in the price facing farmers in former subsidising countries and no change for others.
- As a result, farmers in former subsidising countries produce less. The change in production in each country represents a movement along its supply curve from Q to Q' and can be calculated using the relative change in price ($p^w/(p^w+s)$) and the price elasticity of supply (ϵ_s).
- Production changes in individual countries can then be added together to give the total leftwards movement of the world supply curve (from Q to Q'), reflecting the fact that some production in former subsidising countries is effectively withdrawn from the market.
- A new market equilibrium is then established at the intersection of the new world supply curve and the world demand curve. Equilibrium price and quantity can be found by solving the supply and demand equations simultaneously, using ϵ_s and the price elasticity of demand (ϵ_d).
- At the new equilibrium, production is partially redistributed from former subsidisers to non-subsidising countries. The latter group benefits from increased export earnings due to two factors: firstly, the quantity that would have been exported previously is still exported, but at a higher price; secondly, an additional quantity is exported at the new (higher) price. (In calculating the net figure, an adjustment is made for the cost of imported inputs and it is assumed that local consumption and exports of other products are unaffected.)

Real-world data on prices, subsidies and production are readily available. Supply and demand elasticities are also available, at least for some countries (a problem to which I return below). It is therefore quite straightforward to translate the model's graphical and algebraic relations into numerical ones by plugging in the relevant data and making the necessary calculations. It

was by following just such a procedure that Goreux (2004b) substantiated the claims that, firstly, the world price would have been on average 12% higher over the period 1997-2001 if subsidies had been eliminated and, secondly, that as a result of artificially depressed prices the WCA countries had lost \$250m in export earnings.

As standard as this approach is, its operationalisation nonetheless involves significant limitations that should be made as clear as possible so that its results can be interpreted in an appropriate way. The most important of these will now be discussed in turn.

1.1.1. The Elasticities Problem

Firstly, the model assumes knowledge of the two price elasticities, ε_s and ε_d . Goreux (2004b) does not estimate them using an econometric model, but rather uses the numbers 0.5 and -0.1 respectively. These are based on ICAC/FAO estimates of 0.47 (ε_s for the USA) and -0.06 (world ε_d) and are applied uniformly to all countries. The assumption of uniformity is, of course, a heroic one and would be most unlikely to be supported by empirical investigation. Indeed, Coleman & Thigpen (1991) found demand elasticities ranging from -0.02 (India) to -0.33 (Korea). The same study reported supply elasticities from 0.07 (North India) to 0.87 (Argentina) in the short run and from 0.4 (Central Africa)² to 1.4 (Argentina) in the long run. (The short run supply elasticity for Central Africa was found to be 0.12). The authors also reported results from a handful of other studies, disclosing a wide range of estimated elasticities (from -0.09 to -0.34 for demand, 0.06 to 0.59 for short run supply and 0.25 to 2.46 for long run supply) varying according to country, time period and econometric methodology. Parameter problems of this sort arise frequently in empirical work and Goreux recognised the importance of dealing with them squarely (Goreux (2004b), pp. 27-28). He therefore conducted a sensitivity analysis using different combinations of elasticities, which predictably

² In that study, Central Africa was defined as an aggregate of 22 countries, including the WCA countries.

resulted in a wide range of estimates of the world price effects of removing subsidies, with increases ranging from 2.9% to 13.4%. Similarly, the change in export earnings varies from \$37m to \$254m. His preferred figures of 12% (price) and \$250m (export earnings) can be seen to be well towards the top of this range.

1.1.2. The Stocks Problem

Secondly, the model is based on a standard market-clearing condition that leaves no room for changes in the levels of stocks. Again, this is recognised in the report (Goreux (2004b), p. 27), but in this case is not dealt with comprehensively. It is true that such an assumption is no doubt defensible for a market in which stocks play only a marginal role. Yet today's cotton market stands in stark contrast with that situation: Figure 2 shows that beginning season stocks have in fact fluctuated considerably over the last 40 years and have indeed averaged around 40% of current world consumption since 1970. Such stocks are obviously large enough to have a significant impact on pricing behaviour in both the short- and the long-run. Indeed, Figure 2 suggests that production and consumption are often mismatched on a one-period basis, meaning that fluctuations in the levels of stocks are a vital aspect of market clearing. Relaxation of the "no stocks" condition might therefore be reasonably expected to impact significantly on the final results.

1.1.3. The Comparative Static and Partial Equilibrium Problems

It will be obvious from the outline above that the Goreux (2004b) model compares two static equilibrium states, without describing the dynamics of moving from one to the other. Similarly, it effectively seals the cotton market off from other sectors of the economy and considers it as an autonomous unit. While both assumptions are common in theoretical work, they may differ from reality in ways that are important in terms of the results produced. For instance, ICAC (2002) notes that the price rise thought to flow from removing subsidies would tend to be eaten away in the medium to long term due to the resulting downward

pressure on world demand and upward pressure on production in efficient producer countries. In light of the fierce competition between cotton and synthetic fibres of which Baffes (2004) provides details, the demand effect could even be of short-term relevance if an incipient rise in the price of cotton were to trigger a rapid and strong response in terms of substitution. Again, what this means in concrete terms is that the cited increase in the world price should perhaps be seen as being well towards the top end of the likely range.

1.1.4. The “Design of Experiment” Problem³

The problems identified in the preceding paragraphs are suggestive of another, more fundamental difficulty with this kind of approach. It is, of course, no problem for a theoretical model that some or all of its hypotheses—such as uniform elasticities, no stocks, *ceteris paribus* and so on—should be demonstrably false. What matters is whether the model itself is nonetheless “useful” in terms of the scientific inquiry of which it is part. One very important way in which a model can be “useful” is by making predictions that can be tested in some way against reality. Yet the predictions of the Goreux (2004b) model as it currently stands are unlikely ever to be directly tested, as they would require rich countries to agree to eliminate their subsidies. Such predictions must, therefore, remain a matter of pure theory. Following this logic through, it could even be said that the study itself is not empirical in any real sense, but is more like a simple “translation” of theory from algebra into numbers.

Behind the Goreux (2004b) model, there is, however, a central theoretical construct that is capable of being tested provided the model is reformulated slightly. Specifically, the model assumes the existence of constant elasticity supply and demand curves in all producing and consuming markets (and the aggregated world market), with uniform elasticities across all

³ The term comes from Haavelmo (1944), whose approach to empirical economics provides the essential basis for the methods used here. I have also drawn on later work in a similar spirit, notably Johansen (1995), Sims (1980) and in particular a draft manuscript by Prof. Katarina Juselius entitled “The Cointegrated VAR Model: Econometric Methodology and Macroeconomic Applications” (<http://www.econ.ku.dk/okokj/>).

countries. Tables 1 and 2 show the results of a simple exercise designed to test the fit between that structure and historical data by assessing its ability to predict changes in prices or production between 1997 and 2001 (on a one-step-ahead basis), given information on the other variable. Subsidy-inclusive prices are used for the EU, US and China, meaning that the “experiment” takes into account historical variations in subsidies policies and should provide some indication of the model’s ability to translate those variations into production and price effects. (ICAC data on prices, subsidies and production are used as they mirror very closely the data used in the original report.)

The striking fact that emerges from Tables 1 and 2 is that the reformulated model’s predictive ability is very poor. In terms of price, the average forecasting error (mean absolute percentage error, MAPE) over the five years in question was around 14.5%—greater than the 12% average price increase found to result from the elimination of all subsidies. While production figures are forecast slightly more accurately—a MAPE of 11.5%—the error involved is still large enough to raise serious questions over the model’s “fit” with observed market reality.

Of course, this experiment is far from being a perfect test of the Goreux (2004b) model. As noted above, the model is partial equilibrium and comparative static in nature, whereas the tests conducted are inevitably “contaminated” by general equilibrium and dynamic factors. Does this invalidate the tests? Not necessarily. An alternative interpretation would highlight the fact that the cotton market is neither hermetically sealed from the broader economy nor static in character, meaning that such restrictions should only be imposed in empirical work if they can be justified on the basis of observation. In addition to the obvious problems of elasticities and market clearing, it may be that the factors that Goreux (2004b) suggests would tend to erode an increase in the world price over the medium to long term might in fact act much more quickly and strongly than anticipated. In any case, the experiment can be taken as

demonstrating that some of the assumptions underlying the model appear to prevent it from capturing important aspects of the reality it seeks to describe.

1.1.5. The “Bootstraps” Problem

It is important to keep the Goreux (2004b) model in perspective. As noted above, it is essentially theoretical in character. The numbers it produces—a 12% price increase, \$250m in lost export earnings—are simple quantifications of the theory. In an empirical sense, the model’s numbers cannot, therefore, be used to establish a causal connection *in fact* between higher subsidies and lower world prices. That connection is imposed *a priori* by the model’s theoretical structure, to which the numbers themselves add nothing. It is important to distinguish that approach to modelling from alternative strategies that examine causation in a way that is explicitly empirical (i.e., based directly observation and inference).

When Goreux’s (2004b) results are cited in the public debate as some sort of empirical evidence that rich country subsidies lead to lower world prices, a “bootstraps” argument is involved: that is, the model’s numbers are used as proof of the existence of a causal link which the model itself assumes as a matter of *a priori* theory. Such misappropriation of those results is unfortunately as widespread as it is misleading.

1.2. The ICAC Models

1.2.1. Overview of the Price Model⁴

The ICAC Price Model is essentially a forecasting tool, consisting of a single equation regression relating prices, China’s export market share and stock-to-use ratios (Valderrama, 2000). Only basic diagnostic statistics are presented in Valderrama (2000), but they appear at first glance to demonstrate a reasonable fit between the model and the data: R^2 is just under 0.8, non-constant regressors are statistically significant (based on t-tests) and the regression

⁴ My sincere thanks to Gérald Estur for providing details of these models.

standard error is quite reasonable (five or seven cents, depending on the version used). More importantly, the model's predictive ability seems reasonably good: a one-step-ahead out of sample forecast for the 1999-2000 crop year is only two cents wide of the observed value, while out of sample forecasts for 1998-1999 and 1999-2000 are accurate to within 5 cents. In summary, the model appears to be a useful tool for describing the evolution of world prices.

1.2.2. Overview of the World Textile Demand Model

In a similar vein, the second ICAC model (ICAC, undated) consists of a set of equations used to forecast world textile consumption, covering both cotton and competing natural and synthetic fibres. The structure of the model allows total demand to be decomposed into cotton and non-cotton market shares across three main country groups (industrial countries, developing countries and former Eastern Bloc countries). As was the case for the Price Model, basic regression results seem reasonably sound with high R^2 statistics and a good number of statistically significant regressors. Moreover, the Demand Model's forecasting performance seems very strong, both for cotton and other fibres: its MAPE for one-step-ahead forecasts of consumption over the period 1989-2002 was only around 1.6% in both cases. Again, the model is "useful" in the sense of making quite accurate predictions for the future evolution of the variables under consideration.

1.2.3. Using the Models to Assess the Impact of Subsidies

Given that the two ICAC models provide good price and demand forecasts without making any direct reference to subsidies data, it is somewhat surprising that they have been used to point to the existence of very strong subsidies effects on the world market. Why is it that when zero values are imposed on a variable that has fluctuated significantly in the past but which is excluded from a seemingly useful forecasting model, it turns out that the variables being forecast undergo very substantial changes? For example, the ICAC models have been

cited as suggesting that the world price could increase by as much as 70% if subsidies were to be removed (ICAC (2003), referring to the 2001-2002 crop year).

In light of this disparity, it is important to analyse the process that was used to derive the results in question. From the outline given in ICAC (2002), it would appear that the procedure went essentially as follows:

- Given that US subsidies in the 1994-1995 crop year were very low, actual US production in that year was taken as a proxy for US production at world prices (i.e., without subsidies).
- Using actual world price data for the following years, US production was extrapolated from that baseline figure using a supply elasticity of 0.47 (and, presumably, a simple logarithmic relation between price and quantity changes).
- The extrapolated production figures were then inserted in place of actual production data in the ICAC Price Model and world price forecasts were made on that basis.
- The Price Model's forecasts were fed into the ICAC World Textile Demand Model in order to assess the likely demand effects of changed prices.
- The resulting demand effects were then plugged back into the Price Model to give an indication of the extent to which falling demand would cancel out price increases in the medium to long term.
- The price impact of removing subsidies was then further attenuated by taking into account production increases in non-subsidising countries. It was suggested that half of the production decline would, over time, be offset in this way.
- Once all of these competing forces had been taken into account, a consolidated price impact figure was produced.

- The exercise was then repeated for other subsidising countries, assuming identical production responses to those hypothesised for the US case.

Although this approach is considerably more complex than that of Goreux (2004b), it is nonetheless subject to a number of similar criticisms.

1.2.4. The “Design of Experiment” Problem

Taken separately, the two ICAC models appear to have been the subject of reasonable experiments designed to test their “fit” with observed reality by considering the accuracy of the forecasts they produce. But this is by no means the end of the story, as in this case the models have been used *together* to produce joint forecasts of prices, demand and production. No experiment has been designed to test the new, joint model’s performance in terms of the considerably more difficult task it has been asked to perform. Similarly, there have been no tests of the hypotheses that have been used to link the two models together, or to produce the first series of simulated production data. It is no doubt possible to envisage experiments that could be used to test each of these aspects, but in the absence of published results it is impossible to know the outcome of such tests in advance.

1.2.5. The Linkage Problem

Joint use of the Price and Demand models gives rise to two more problems of a technical nature, but which are likely to have a major impact on the results obtained as well as on their interpretation. Firstly, the two models are effectively treated as forming an interlinked system of equations despite the fact that each was estimated by ordinary least squares (OLS) without reference to the other and hence could easily be subject to simultaneity bias. It would have been more appropriate to estimate the two models together from the start, using a suitable systems estimator. Secondly, and related to this point, the mechanical transfer of estimates from one model to another ignores the forecasting error inherent in each step. For example, the production estimates made on the basis of a constant elasticity function applied to 1994

data are, like any other statistical projections, subject to a certain amount of error. However, when they are plugged into the Price Model, they are treated as data rather than forecasts and thus no allowance is made for such uncertainty. The more this process is repeated, the larger the overall error is likely to be. In presenting results obtained in this way without including appropriate confidence bounds, the ICAC methodology can easily be misinterpreted as making predictions that are subject to only negligible error when, in reality, there is no indication as to how large that error might in fact be.

1.2.6. The “Bootstraps” Problem

Finally, the ICAC approach is subject to similar reserves concerning causation as were expressed in relation to Goreux (2004b). The modelling approach imposes a causal order running from subsidies to production to prices to demand, then back to prices, production and prices again. In theoretical terms, this is just one possible (untested) system of causal linkage amongst many. The results it produces must, therefore, be taken for what they are. As was the case for Goreux (2004b), it is difficult to see how quantitative results that flow directly from a causal structure imposed *a priori* can later be cited as evidence that the causal chain in question actually exists as a matter of empirical fact. That ICAC’s results have effectively been assigned such a role in the public debate is again indicative of serious and unfortunate errors of interpretation.

1.3. Results from Other Models

Table 3 puts the Goreux and ICAC results in context, by comparing them with a number of other studies that have sought to assess the effects of removing cotton subsidies. It is not possible to directly compare world price effects across models, as liberalisation scenarios differ significantly. However, the range of reported price results is nonetheless remarkable, extending from +2% to +70%. The Goreux (2004b) and ICAC estimates are certainly towards

the high end of this range, an important fact to note given their relative importance in the political debate.

It is not possible here to enter into a detailed criticism of the methodologies at the base of each of the studies cited in Table 3.⁵ However, it should be clear that many of the same comments apply. That is, the results are essentially theoretical, rather than empirical, in that they reflect the assumptions of the modellers as to economic structure, rather than the interrelationships that emerge over time from the data themselves. This is just as true of a complex computable general equilibrium model as it is of a simple partial equilibrium simulation. It is important to emphasise that there does not as yet appear to be any study that has attempted to analyse the effects of subsidies on the cotton market in a purely empirical way. It is with this in mind that the next section outlines an alternative strategy for approaching the subsidies issue.

2. An Alternative Modelling Strategy

Following Sims (1980), and mirroring techniques widely used in macroeconometrics and macroeconomic policy analysis (Sims, 1986), I propose approaching the subsidies question in two stages. Firstly, I will attempt to build a subsidy-inclusive model of the world cotton market that is meaningful from a statistical point of view, but which is as devoid as possible of *a priori* restrictions derived from economic theory. I will then test the model against historical data to ensure that it is a reasonably close fit. (In other words, I will test whether it is “useful” to impose a theoretical fiction, namely that the observed stream of data was produced by the statistical processes included in the model; cf. Haavelmo (1944).) Once the model’s ability to reproduce observed data has been demonstrated, I will move to the second

⁵ In relation to the work by Reeves et al (2001), such an assessment cannot even be contemplated, as the models used are regarded by their creators as “commercial-in-confidence” (e-mail communication from George Reeves to the author, dated 20 January 2004). These results can therefore never be subject to independent review.

stage of this strategy, in which the model is used as a statistical “test-bed” upon which economically meaningful experiments can be performed. Hypotheses of interest—such as “subsidies lead to lower prices”—will be translated into mathematical terms and tested using statistical methods. It is on the basis of the results from these hypothesis tests that conclusions will be drawn as to the impact that subsidies have in fact had on the cotton market in the past and, by extension, the likely effects of their elimination or reduction.

To implement the first step in this strategy, a very general statistical structure is needed to translate the data into a mathematically convenient form. As has often been the case in other contexts, a simple vector autoregression (VAR) was found to do the job well.

The basic VAR model is not (initially at least) motivated by any particular economic theory and does not have any obvious economic interpretation, with each variable simply modelled as a function of its own past values and those of other endogenous variables. All variables are initially taken as endogenous, meaning that no causal order is imposed on the data *a priori*. In more concrete terms, what the VAR model gives us in this case is just a shorthand mathematical representation of a group of variables thought to describe the workings of the cotton market, but without the prior imposition of ideas from theoretical economics that are believed to govern the variables’ inter-relationships.

3. Building a VAR Model of the World Cotton Market

As a starting point for investigating the impact of subsidies on the world cotton market, it was decided to focus exclusively on the effects of US subsidies. This choice was made for a number of reasons. Firstly, a long and uniform data series is available to approximately track US subsidy policy over a sufficient period to make VAR modelling viable. Secondly, US subsidies have been presented politically as the primary culprits behind the subsidy-induced decline of world prices (e.g., Oxfam, 2002). Finally, previous quantitative work has suggested that US subsidies *on their own* have significant world price effects (see Table 3).

3.1. Data

The dataset consists of annual⁶ observations between 1965 and 2002 covering prices, world production, consumption and stocks, and US subsidies. The series are defined in the following terms (see Table 13 for further details on sources and units and Figures 2-5 for graphs in real and nominal terms where appropriate):⁷

- p_t : the nominal world price for cotton, proxied by the IMF's "Liverpool Price". This is obviously a simplification, as numerous different types and grades of cotton are traded on the world market and such complexity cannot hope to be captured in a single price. However, such a simplification is necessary to make the model tractable and also reflects industry practice, in which the Cotlook "A" Index—which differs only marginally from the IMF series—is used as a broad indicator of world market prices.
- q_t : total world production of cotton, based on data from the US Department of Agriculture (USDA).
- x_t : total world consumption of cotton, based on data from the USDA.
- z_t : the change in the level of total world stocks of cotton over period t , based on data from the USDA.
- s_t^{US} : US subsidies to cotton producers per pound of production, proxied by US (federal) government direct payments to the cotton sector (in nominal terms). This is by no means an ideal proxy, as it includes neither federal government loan payments, nor direct and

⁶ Years refer to US crop years (August to July).

⁷ Tables 4 and 5 show the results of pre-testing and immediately prompt doubts as to whether all series are stationary, in particular in light of the fact that different sample periods and frequencies of related series have been treated as I(1) in other work, such as Baffes & Ajwad (1998). Although the weight of evidence seems to point towards stationarity with a structural break due to the oil shock appearing to some tests like a unit root, the matter is still open to question and is the subject of ongoing research.

indirect payments by individual states. As a result, it probably understates the true level of subsidisation. However, it would appear to be the best series available, and indeed closely tracks data for the late 1990s presented in ICAC (2002) and Goreux (2004b). The series also has the important advantage of covering a much longer historical period than other possible proxies, such as the WTO's product-specific Aggregate Measure of Support (from 1995 only).

For modelling purposes, all variables are expressed in logarithms.

3.2. Model Form

The general VAR(j) model in matrix form is simply:

$$Y_t = \mu + \sum_{i=1}^j \beta_i Y_{t-i} \quad (1a)$$

where Y is the vector of endogenous variables (initially all data series), μ is a vector containing deterministic terms and β is the coefficient matrix. In expanded form, this becomes:

$$\left. \begin{aligned} p_t &= \mu^0 + \mu^d DU + \mu^t t + \sum_{i=1}^j \beta_i^p p_{t-i} + \sum_{i=1}^j \beta_i^q q_{t-i} + \sum_{i=1}^j \beta_i^x x_{t-i} + \sum_{i=1}^j \beta_i^z z_{t-i} + \sum_{i=1}^j \beta_i^s s_{t-i} + \varepsilon_t^p \\ q_t &= \mu^0 + \mu^d DU + \mu^t t + \sum_{i=1}^j \beta_i^p p_{t-i} + \sum_{i=1}^j \beta_i^q q_{t-i} + \sum_{i=1}^j \beta_i^x x_{t-i} + \sum_{i=1}^j \beta_i^z z_{t-i} + \sum_{i=1}^j \beta_i^s s_{t-i} + \varepsilon_t^q \\ x_t &= \mu^0 + \mu^d DU + \mu^t t + \sum_{i=1}^j \beta_i^p p_{t-i} + \sum_{i=1}^j \beta_i^q q_{t-i} + \sum_{i=1}^j \beta_i^x x_{t-i} + \sum_{i=1}^j \beta_i^z z_{t-i} + \sum_{i=1}^j \beta_i^s s_{t-i} + \varepsilon_t^x \\ z_t &= \mu^0 + \mu^d DU + \mu^t t + \sum_{i=1}^j \beta_i^p p_{t-i} + \sum_{i=1}^j \beta_i^q q_{t-i} + \sum_{i=1}^j \beta_i^x x_{t-i} + \sum_{i=1}^j \beta_i^z z_{t-i} + \sum_{i=1}^j \beta_i^s s_{t-i} + \varepsilon_t^z \\ s_t &= \mu^0 + \mu^d DU + \mu^t t + \sum_{i=1}^j \beta_i^p p_{t-i} + \sum_{i=1}^j \beta_i^q q_{t-i} + \sum_{i=1}^j \beta_i^x x_{t-i} + \sum_{i=1}^j \beta_i^z z_{t-i} + \sum_{i=1}^j \beta_i^s s_{t-i} + \varepsilon_t^s \end{aligned} \right\} (1b)$$

Three deterministic terms are included in the model, namely a constant, a dummy variable that is zero up to 1973 and 1 thereafter, and a time trend. The dummy is included to take account of the price hike that occurred following the 1970s oil shock, when production costs

of synthetic fibres increased substantially and as a result, there was a marked demand switch towards cotton that in turn pushed cotton prices up. The time trend is designed to take account of strong trending behaviour observed in the production and consumption series as well as to improve model stability (as those two series appear to be trend stationary).

In its basic form, the VAR model can be estimated by simple equation-by-equation OLS regression (Hamilton, 1994). Standard Wald tests are commonly used for hypothesis testing, though given the very small sample under consideration here it is more appropriate to use the modified likelihood ratio test suggested by Sims (1980):

$$\lambda = (T - c)(\log|\Sigma_R| - \log|\Sigma_U|) \sim \chi^2_{df=\text{no. of restrictions}} \quad (2)$$

3.3. Estimation, Diagnostic Tests and Forecast Evaluations

Basic diagnostic results for a VAR(4) model are presented in Tables 6 (equation by equation tests) and 7 (systems tests).⁸ A lag length of four (the maximum considered, due to limited data) was chosen as it was found to minimise the Akaike Information Criterion (AIC). Model residuals appear to be normally distributed, though on the system level the result is sensitive to methodology. Recursive estimation-based tests indicate that the model parameters are reasonably stable over time, although there is some weak evidence of structural change late in the sample.⁹ Statistical tests also suggest that some residual autocorrelation appears to be present—it is persistent to changes in the number of lags used—but visual inspection of the correlograms suggests that the problem is perhaps not as serious as it at first appears from the LM test statistics. At any rate, individual equation Q-statistics for 12 lags are only 5% significant in one case, again suggesting that the problem is perhaps not an overly serious one.

⁸ All calculations were performed using eViews 4.1. Estimated coefficients are not reproduced here as they are of little interest given the modelling strategy used. However, a more detailed Statistical Appendix containing these and other results excluded from the present paper is available from the author on request.

⁹ Tests performed included inspection of recursive residuals, CUSUM, CUSUM of squares, forecast tests and recursive coefficient estimates.

Nonetheless, attention will need to be given in future work to finding ways of better dealing with this issue.

In light of these results, the model seems on the whole to be tolerably well specified from a statistical point of view. It is also very good at reproducing historical data in terms of one-step-ahead, within-sample forecasts: its mean absolute error (MAE) for price forecasts is around 3.6 cents over the 1969-2001 period, which equates to a MAPE of just under 6%. Comparable tests for quantities produced and consumed give even more impressive results (see Table 8), though forecasts of stock movements and of US subsidies themselves are considerably less accurate. On a dynamic, within-sample basis from 1997-2001 (Table 9), results are still solid though obviously involve a larger margin for error than in the static case (the price equation MAPE increases to just over 6%, for example).

Given these strong within-sample results, it is a little disappointing that the model's out-of-sample forecasting ability seems somewhat weaker (see Table 10). When data through 2001 are used to forecast 2002 values, the predicted price is quite wide of the mark (+8.7 cents or just over 15%). Although production and consumption forecasts are considerably better, the model's out-of-sample forecasting ability remains noticeably inferior to its within-sample performance. Taken with the evidence of late-sample parameter instability referred to earlier, it may be that this effect is due in part to weak structural change towards the end of the period.

To put the model's performance in perspective, it is useful to benchmark it against the other models that have been reviewed earlier in this paper.¹⁰ A comparison of Tables 1-2 and 8 shows that the VAR model's forecasts outperform those of the reformulated Goreux (2004b)

¹⁰ As a baseline, we can also consider the model's performance compared with that of a simple random walk. For prices (1969-2001), the random walk's MAE is nearly 12 cents and its MAPE is 18%. Its prediction for 2002 comes with an error of nearly 14 cents. The VAR model therefore clearly outperforms the random walk.

model. Moreover, its within-sample performance would appear to compare quite favourably with what is known of the ICAC Price and Demand Models.¹¹ The only black spot, as just noted, is its somewhat poorer performance out-of-sample, namely for the 2002 crop year. While that result is of concern and the subject of continuing investigation, it should not obscure the fact that in the main the VAR model appears to reproduce observed reality reasonably well.

3.4. Making the Model “Talk”

Now that the model has been found to be appropriately specified, it can be used to investigate the dynamic interactions between the different variables using Impulse Response Function (IRF) plots and Forecast Error Variance Decompositions (FEVDs). Following Bernanke (1986), a structural factorisation was used to impose some minimal identifying restrictions on the model’s errors. Treating the observed errors in system (1b) as functions of uncorrelated structural disturbances (v_t), we have $A\varepsilon_t = \Gamma v_t$, where:

$$A = \begin{bmatrix} 1 & \alpha_{12} & \alpha_{13} & \alpha_{14} & \alpha_{15} \\ 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 \\ 0 & \alpha_{42} & \alpha_{43} & 1 & 0 \\ 0 & \alpha_{52} & \alpha_{53} & \alpha_{54} & 1 \end{bmatrix}; \Gamma = \begin{bmatrix} \gamma_{11} & 0 & 0 & 0 & 0 \\ 0 & \gamma_{22} & 0 & 0 & 0 \\ 0 & 0 & \gamma_{33} & 0 & 0 \\ 0 & 0 & 0 & \gamma_{44} & 0 \\ 0 & 0 & 0 & 0 & \gamma_{55} \end{bmatrix} \quad (3)$$

In summary, this identifying scheme allows prices to respond contemporaneously to unexpected changes in all other variables, as is appropriate given their role in re-establishing market equilibrium following shocks. Stocks can adjust immediately to production and consumption shocks—again, this is inherent in their role in the market—while subsidies can respond to unexpected movements in all three of those variables. Production and

¹¹ Using similar performance measures, it also stands up well against the forecasting model in Kaltsas (2000). On the basis of root mean square percentage errors (RMSPE), it holds its own against the far more complicated structural models of Coleman & Thigpen (1991) especially since, as those authors acknowledge, RMSPEs for their models are artificially deflated due to the extensive use of dummy variables to exclude outlier observations.

consumption, on the other hand, are assumed to be directly subject to structural shocks only. The fact that production is so constrained can be justified by recalling that what matters to farmers when making production decisions is the *expected* level of subsidies, as opposed to the unforeseeable shocks under consideration here. But in any case, allowing subsidies to enter the production error term directly does not materially alter the results reported here.

All estimated coefficients from maximum likelihood estimation of A and Γ are significant at the 10% level—all but two are 1% significant—and a likelihood ratio test for overidentification does not reject the null hypothesis that the restrictions are valid (prob. value = 0.6955). Hence, although the identification scheme is but one amongst many that could be considered justifiable in theoretical terms, it has the distinct advantage of not being rejected by the data.

Figure 6 shows cumulative IRF plots (i.e., the accumulated response of each variable to a one period shock to another variable) based on this factorisation. It is apparent that over time, prices respond to unforeseen changes in production and consumption in the way that basic theory would suggest, falling in response to the latter and rising in response to the former. Similarly, production increases following a price shock. More interesting is the consumption response: after initially falling in response to a price shock, it then in fact increases markedly. The reason for this apparently counter-intuitive result would appear to be that the price shock is associated with rising production which then pushes the price down and tends to stimulate consumption.

In terms of the impact of subsidies on the world market, the IRF plots show that unexpected policy changes have a noticeable impact over time not only on production, but also on consumption and stocks. Indeed, consumption appears to increase very strongly in response to a subsidies shock, though its reaction is lagged by three years or so. The world price initially rises marginally, then falls slightly, before apparently reaching a durably higher level than

prior to the shock. This suggests a complex dynamic relationship between subsidies, production, consumption and price that would appear to defy the simple predictions of standard partial equilibrium theory. Although further work is obviously necessary to get to the bottom of that relationship, one possible mechanism that emerges from the IRFs is that the increase in production following a subsidies shock is initially absorbed by larger stocks. This, combined with a significantly lagged consumption response, dampens the expected price effects. Over time, however, stocks are released and prices are pushed down, all of which tends to stimulate consumption. In the end, the upwards pull of consumption could conceivably overcome the downward force exerted by production and thereby lead to an overall price rise, accompanied by a durably higher level of stocks following reaccumulation. One reason why the consumption response might be so significant is that subsidies shocks could foster substitution away from synthetic fibres in the short to medium term, a mechanism that would be in line with the strong level of observed competition between the two industries (Baffes, 2004), but which needs further empirical work before it can be considered as having been firmly substantiated.¹²

Interesting though these effects are, it is also necessary to keep their extent in perspective: Figure 7 shows the same IRF plots with two standard error bands, indicating that in many cases the effects discussed above are of only marginal statistical significance. So it is in fact probable that a shock to US subsidies has only a negligible impact on price levels, with production and consumption effects seemingly cancelling each other out. On the other hand, the oscillatory effects observed could perhaps contribute to price volatility even if their effect on levels is minimal.

¹² Baffes (2003) cites the Director-General of the International Rayon and Synthetic Fibres Committee: “recent increases in cotton subsidies have rigged the market even more dramatically in favour of cotton, depressing demand for every substitute product”.

This scepticism as to the strength of US subsidies' impact is confirmed by the FEVDs presented in Figure 8. Subsidies shocks are found to account for only around 10% of the forecasting errors affecting prices and production and approximately 20% for consumption. The largest single component of price forecast error is in fact consumption shocks, suggesting that the demand side of the market—and notably the strong potential for substitution following price changes—is of particular importance in price determination.

3.4.1. Do US Subsidies “Cause” Low Prices?

Having outlined the system's dynamics in terms of each variable's response to unforeseen shocks to the other variables, I will now examine the impact of subsidies in more historical terms. Specifically, the hypothesis that US subsidies “cause” low prices will be tested by focussing on its narrow statistical sense (Granger causality) to the effect that past values of the former constitute a valuable piece of information for forecasting the current value of the latter (Hamilton, 1994). In formal terms, Sims' (1980) likelihood ratio test is used to test the joint null hypothesis ${}_p\beta_i^s = 0$ for $\forall i = 1 \dots 4$.

The first line of Table 11 shows that the test does not reject the null hypothesis, meaning that historical levels of US subsidies are not particularly useful in making forecasts of current world prices. In concrete terms, this means that it is quite possible to make reasonable forecasts of world cotton prices without the slightest indication as to the level of US subsidies in the past—as, indeed, the mainstream price models reviewed above all do.¹³

However, the full story is not quite that simple. Table 11 suggests that although past values of subsidies may not be of much direct relevance for current prices, they are more important for

¹³ To make the point abundantly clear, a VAR model excluding subsidies entirely was also estimated and it was found to perform very well against the same criteria applied to the original model. Indeed, it exhibited considerably less evidence of residual autocorrelation than did the original model, suggesting that no excluded variable—such as subsidies—caused systematic correlation amongst the model's error terms. This is a further piece of evidence tending to suggest that the impact of US subsidies is quite limited.

forecasting production, consumption and movements in stocks. It is therefore of interest to perform a more general test to see whether subsidies play an important role in forecasting on a system-wide basis. In formal terms, the joint null hypothesis is ${}_k\beta_i^s = 0$ for $\forall i = 1..4$ and $\forall k = p, q, x, z$. The likelihood ratio test statistic is 24.823372, which with a prob. value of 0.073000 is only marginally statistically significant (i.e., at the 10% level, but not at the 5% or 1% levels). This test, like the others, tends to support the impression that in historical terms, US subsidies are not a very important determinant of the current state of the system.

It is also of interest to test whether subsidies “belong” in the VAR at all, that is whether they appear to be a genuinely endogenous variable (as has been assumed). The null hypothesis for block exogeneity of subsidies is simply ${}_s\beta_i^k = 0$ for $\forall i = 1..4$ and $\forall k = p, q, x, z$. Table 11 shows that, once again, the null cannot be rejected. This suggests that the data might allow a respecification of the model, treating subsidies as exogenous:

$$\left. \begin{aligned} p_t &= {}_p\mu^0 + {}_p\mu^d DU + {}_p\mu^t t + \sum_{i=1}^j {}_p\beta_i^p p_{t-i} + \sum_{i=1}^j {}_p\beta_i^q q_{t-i} + \sum_{i=1}^j {}_p\beta_i^x x_{t-i} + \sum_{i=1}^j {}_p\beta_i^z z_{t-i} + {}_p\beta^s s_t + \varepsilon_t^p \\ q_t &= {}_q\mu^0 + {}_q\mu^d DU + {}_q\mu^t t + \sum_{i=1}^j {}_q\beta_i^p p_{t-i} + \sum_{i=1}^j {}_q\beta_i^q q_{t-i} + \sum_{i=1}^j {}_q\beta_i^x x_{t-i} + \sum_{i=1}^j {}_q\beta_i^z z_{t-i} + {}_q\beta^s s_t + \varepsilon_t^q \\ x_t &= {}_x\mu^0 + {}_x\mu^d DU + {}_x\mu^t t + \sum_{i=1}^j {}_x\beta_i^p p_{t-i} + \sum_{i=1}^j {}_x\beta_i^q q_{t-i} + \sum_{i=1}^j {}_x\beta_i^x x_{t-i} + \sum_{i=1}^j {}_x\beta_i^z z_{t-i} + {}_x\beta^s s_t + \varepsilon_t^x \\ z_t &= {}_z\mu^0 + {}_z\mu^d DU + {}_z\mu^t t + \sum_{i=1}^j {}_z\beta_i^p p_{t-i} + \sum_{i=1}^j {}_z\beta_i^q q_{t-i} + \sum_{i=1}^j {}_z\beta_i^x x_{t-i} + \sum_{i=1}^j {}_z\beta_i^z z_{t-i} + {}_z\beta^s s_t + \varepsilon_t^z \end{aligned} \right\} \quad (4)$$

This adapted model can be used to test for the instantaneous impact of subsidies on prices by imposing the null hypothesis that ${}_p\beta^s = 0$. As Table 12 shows, the data do not reject the null for the price equation, but do reject it strongly for the production, consumption and stocks equations. A system-wide test using the null ${}_s\beta^k = 0$ for $\forall k = p, q, x, z$ is only narrowly rejected at the 5% level (prob. value = 0.045697). Putting these results together broadly

confirms the view taken above, namely that while subsidies appear to impact significantly both on production and consumption, their overall effect on the market—and more particularly, on the world market price—is surprisingly limited.

3.4.2. Simulating the Effects of a Reduction in US Subsidies

As a final exercise, the VAR model with exogenous subsidies is used to perform a counterfactual “what if?” simulation (cf. Myers et al, 1990), in which a percentage reduction is applied to per pound subsidies data from 1997 to 2001 and a stochastic dynamic simulation is run. Prices, production, consumption and stocks are jointly simulated on the basis of the altered subsidies data. Such an exercise assumes that the VAR model as estimated would continue to describe the data accurately, even if the level of subsidisation were to be changed. This seems a reasonable hypothesis at least for small changes, and is supported by the fact that the model has been shown to be good at reproducing historical data in the presence of fluctuating subsidies. The market as observed from the outside appears “as if” it had been produced by the statistical processes summarised in the VAR, meaning that when a small amount of counterfactual “data” is added we can expect that the VAR should reproduce the market reality that we would be likely to observe if the “data” were in fact real.

Figures 9 to 11 contain the results of three such counterfactuals in which subsidies are reduced by 10%, 50% and 90% respectively. (“Scenario One” refers to the simulation based on reduced subsidies, while “Baseline” uses actual subsidies data and is the dynamic simulation reported above.) A 10% reduction in US subsidies is seen to have essentially no impact on the market, with simulated prices and quantities almost indistinguishable from the baseline. While the 50% and 90% reductions have some impact on the simulated means of the price and quantity variables—in both cases, the world price appears to increase—it is important not to overstate these impacts given the very significant degree of uncertainty that is inherent in a multivariate, medium-horizon forecasting exercise such as this one. This

uncertainty is reflected in the fact that the 95% confidence intervals for the reduced subsidy simulations are very wide. Indeed, they are so wide that it would be dangerous to conclude that the hypothesised effects are statistically significant at all. In short, simulation results do not provide much support for the view that cutting US subsidies, even drastically, is likely to result in very different aggregate market conditions; rather, they highlight the complexity of the different forces at play and the difficulty in making meaningful forecasts that attempt to take that complexity into account.

4. Conclusion and Directions for Future Research

I would suggest that the VAR-based results reported here give cause for a serious reassessment of the effects that can be ascribed *in fact* to US subsidies. In contrast to previous work that has attempted to quantify those effects using primarily theoretical methods, the empirical approach adopted here suggests that subsidies are not an overly important determinant of world prices' expected behaviour, either in historical or contemporaneous terms. Moreover, unexpected changes in US subsidy policy have been shown to have the potential to produce such apparently "perverse" effects as a rise in prices over the medium to long term. Simulation results indicate that although a reduction in US subsidies could in principle lead to world prices that are to some degree higher, the effect is not significant in a statistical sense; in other words, it is perfectly possible that even a 90% reduction in subsidies could leave world prices unchanged.

There is still considerable work to be done to confirm that these results are robust. Ongoing research in that direction is concentrating on the following points:

- The issue of possible non-stationarity of some of the series needs to be addressed due to its implications for hypothesis testing (see Sims, Stock & Watson, 1990). Given the likely mix of stationary and non-stationary variables, the lag-augmented VAR approach of Toda

& Yamamoto (1995) could be a useful way of confirming the results found using a standard VAR model (see also Yamada & Toda, 1998).

- However, model stability needs to be addressed, as early efforts to use lag-augmented VARs for hypothesis testing have run into the difficulty that the augmented models have characteristic roots in excess of unity (whereas the base models do not).
- The related question of possible late-sample parameter shifts also needs to be dealt with in more detail.
- Given the relatively small amount of data and large number of coefficients to be estimated in each model, degrees of freedom represent a serious constraint for hypothesis testing. One possible way to address this is to “prune” the VAR model so as to exclude the regressors of least explanatory power in each equation, applying Seemingly Unrelated Regression (SUR) to estimate the resulting unbalanced system (Hamilton, 1994) and using AIC or another information criterion to select the final model. Alternatively, a Bayesian prior could be imposed on the lag structure, as originally suggested in Sims (1980).
- Although preliminary tests using different model structures and variables have produced basically similar results, further work is necessary to ensure that the results are reasonably robust to the choice of statistical model.

Leaving these questions to one side for the moment, it nonetheless seems clear that the results in the present paper are quite contrary to expectations formed in terms of the “conventional wisdom” that has grown up around the cotton subsidies debate at the WTO (i.e., US subsidies push world prices down, thereby making African farmers poorer). Indeed, at the outset of this exercise I anticipated finding strong evidence that subsidies “cause” prices, either in historical (Granger) or instantaneous impact terms, and correspondingly that removing US subsidies leads to unambiguously and significantly higher prices. That the data appear inconsistent with

these expectations—or at the very least, consistent with contrary ones—must raise serious doubts as to the empirical applicability of previous quantification work in this area. A healthy dose of scepticism, combined with additional explicitly empirical research, seems called for.

The methodological and factual issues that have been considered in the present paper suggest a number of research questions that could be fruitfully investigated, perhaps by adapting the methods used here:

- The results found here suggest that subsidies may play some role in heightening world price volatility. It would be useful to test this hypothesis more explicitly using a GARCH-type framework in which volatility is modelled explicitly.
- The present model does not explicitly take into account substitution between cotton and synthetic fibres. It would be interesting to include the relative price of cotton in terms of synthetics in the empirical framework. This would allow us to see the extent to which the price effects of changes in subsidies policy might be nullified due to strong demand switching. Such an effect could well be behind some of the results reported here and is deserving of closer analysis.
- It would also be profitable to investigate the US market in detail, in order to better understand the dynamics of the subsidies-quantities-prices relationship that is at the core of the issue in terms of the world market. This would provide a useful point of comparison with previous studies that have modelled cotton production using single or simultaneous equation techniques.
- A crucial aspect of the WCA countries' ability to take advantage of any change in US subsidies policy is their farmers' freedom to respond to any resulting price signals. However, Baffes (2004) points out that in at least some of those countries the state is still heavily involved in the cotton marketing chain, sometimes to the exclusion of private

actors. As a result, producer and world prices may not move together and there is a question as to the degree to which farmers would actually respond to any change in world prices. A VAR model relating producer and world prices with data on quantities could be used to investigate these dynamics. Such a model would also enable useful comparison with estimated price elasticities from previous work, such as Coleman & Thigpen (1991).

- Related to that question is the issue of the extent to which world cotton prices in fact impact on GDP in the WCA countries. Structural VAR models have frequently been used to analyse the determinants of GDP shocks and could easily be adapted to examine that question, on which there does not as yet appear to be any empirical work.
- Baffes (2004) also shows that implicit taxes on cotton farmers in WCA countries are sometimes substantial, with a large gap between world and producer prices. The extent to which rural poverty is a function of that gap, rather than subsidy-induced shocks to world prices, is an open question that could be explored using a similar modelling strategy in which a VAR is used to dynamically relate some poverty measure to world and producer prices.

As well as being of political interest, the WCA countries' sectoral initiative is also, therefore, of considerable economic interest. Rather than stating a series of conclusions that are already firmly grounded in empirical economics, it instead opens up a field of inquiry in which genuinely observation-based research is surprisingly scarce but which has the potential to deliver up significant insights for empirical trade policy.

References

- Baffes, John. 2003. "Cotton and Developing Countries: A Case Study in Policy Incoherence", World Bank Trade Note 10, 10 September 2003, www.worldbank.org.
- Baffes, John. 2004. "Cotton: Market Setting, Policies, Issues, and Facts", World Bank Policy Research Working Paper 3218, February 2004, www.worldbank.org.
- Baffes, John & Mohamed I. Ajwad. 1998. "Detecting Price Linkages: Methodological Issues and an Application to the World Market of Cotton", World Bank Working Paper 1944, Washington, DC: World Bank.
- Benin, Burkina Faso, Chad and Mali. 2003. "Poverty Reduction: Sectoral Initiative in Favour of Cotton", WTO Doc. TN/AG/GEN/4, 16 May 2003.
- Bernanke, Ben S. 1986. "Alternative Explanations of the Money-Income Correlation", NBER Working Paper 1842, February 1986, www.nber.org.
- Brazil. 2002. "US – Subsidies on Upland Cotton, Request for Consultations by Brazil", WTO Doc. WT/DS267/1, 3 October 2002.
- Coleman, Jonathon R. & M. Elton Thigpen. 1991. "An Econometric Model of the World Cotton and Non-Cellulosic Fibers Markets", World Bank Staff Commodity Working Paper No.24, www.worldbank.org.
- FAPRI. 2002. "The Doha Round of the World Trade Organization: Appraising Further Liberalization of Agricultural Markets", Working Paper 02-WP 317, November 2002.
- Goreux, Louis. 2004a. "Cotton After Cancun", Mimeo, January 2004.
- Goreux, Louis. 2004b. "Prejudice Caused by Industrialized Countries Subsidies to Cotton Sectors in Western and Central Africa", 2nd edition, January 2004, www.worldbank.org.
- Greene, William H. *Econometric Analysis*. Fourth edition, New Jersey: Prentice-Hall.
- Haavelmo, Trygve. 1944. "The Probability Approach in Econometrics", *Econometrica*, Vol. 12, Supplement (Jul., 1944), iii-vi & 1-115.
- Hamilton, James D. 1994. *Time Series Analysis*. Princeton, NJ: Princeton University Press.
- ICAC. Undated. "Methodology of World Textile Demand Projections". *Mimeo*.
- ICAC. 2002. *Production and Trade Policies Affecting the Cotton Industry*. Washington, DC: ICAC.
- ICAC. 2003. "Impacts of Government Measures on African Cotton Producers", Statement by the Executive Director to the UNCTAD Trade and Development Board, 13 October 2003, www.icac.org.
- ICAC. 2003a. "Government Support to the Cotton Industry", Mimeo, www.icac.org.
- Johansen, S. 1995. *Likelihood-Based Inference in Cointegrated Vector Autoregressive Models*. Oxford: Oxford University Press.
- Kaltsas, Ioannis K. (2000). "Explaining International Cotton Prices: A Structural Model Approach", *Proceedings of the Beltwide Cotton Conference*, Vol.1, 273-274, www.icac.org.
- Kwiatkowski, Dennis, Peter C.B. Phillips, P. Schmidt & Y. Shin. 1992. "Testing the Null Hypothesis of Stationarity Against the Alternative of a Unit Root", *Journal of Econometrics*, Vol.54, 159-178.

- Myers, Robert J., R. R. Piggott & William G. Tomek. 1990. "Estimating Sources of Fluctuations in the Australian Wool Market: An Application of VAR Methods", *Australian Journal of Agricultural Economics*, Vol.34, No.3, December 1990, 242-262.
- Oxfam. 2002. "Cultivating Poverty: The Impact of US Cotton Subsidies on Africa", Briefing Paper 30, www.oxfam.org.
- Perron, Pierre. 1989. "The Great Crash, the Oil Price Shock, and the Unit Root Hypothesis", *Econometrica*, Vol.57, Issue 6, November 1989, 1361-1401.
- Reeves, George, David Vincent, Derek Quirke & Stephen Wyatt. 2001. *Trade Distortions and Cotton Markets: Implications for Australian Cotton Producers*. Narrabri, NSW: Cotton Research and Development Corporation.
- Sims, Christopher A. 1980. "Macroeconomics and Reality", *Econometrica*, Vol. 48, No. 1, January 1980, 1-48.
- Sims, Christopher A. 1986. "Are Forecasting Models Usable for Policy Analysis?", *Federal Reserve Bank of Minneapolis Quarterly Review*, Winter 1986, 2-16.
- Sims, Christopher A., James H. Stock & Mark W. Watson. 1990. "Inference in Linear Time Series Models with Some Unit Roots", *Econometrica*, Vol.58, No.1, January 1990, 113-144.
- Toda, Hiro Y. & Taku Yamamoto. 1995. "Statistical Inference in Vector Autoregressions with Possibly Integrated Processes", *Journal of Econometrics*, Vol.66, 225-250.
- Valderrama, Carlos. 2000. "A Revision of the ICAC Price Model", Mimeo, July-August 2000, Washington, DC: ICAC.
- Valderrama Becerra, Carlos. 2000. "The World Cotton Market: Prices and Distortions", Mimeo, www.icac.org.
- Yamada, Hiroshi & Toda, Hiro Y. 1998. "Inference in Possibly Integrated Vector Autoregressive Models: Some Finite Sample Evidence", *Journal of Econometrics*, Vol.86, 55-95.

Tables

Table 1: Simulation results based on Goreux (2004b) for endogenous prices.

Country	1998		1999		2000		2001	
	P' Est.	APE	P' Est.	APE	P' Est.	APE	P' Est.	APE
USA	81.62	0.88	79.31	1.94	77.70	12.28	68.06	3.43
China	92.40	7.57	87.30	23.31	69.79	9.60	75.78	46.30
Greece	159.42	11.56	140.10	22.04	115.18	0.02	114.82	13.91
Spain	156.90	6.09	144.42	19.55	124.97	12.73	141.36	20.00
ROW	72.34	22.81	58.75	11.27	52.90	7.51	56.97	36.30

Table 2: Simulation results based on Goreux (2004b) for endogenous production.

Country	1998		1999		2000		2001	
	Q' Est.	APE	Q' Est.	APE	Q' Est.	APE	Q' Est.	APE
USA	4083.32	34.76	3041.86	17.65	3737.53	0.12	3760.90	14.91
China	4634.69	2.97	4588.86	19.84	3796.01	14.12	4599.93	13.54
Greece	343.91	3.67	364.90	16.11	434.85	3.29	426.66	1.92
Spain	116.56	12.08	106.13	19.60	129.77	38.06	95.85	10.42
ROW	11148.69	4.00	10837.84	1.47	10912.30	1.19	11127.61	0.85

Table 3: Summary of Previous Research on the Effects of US Cotton Subsidies¹⁴

Study	Method	World Price Effects of Removing US Subsidies Only	World Price Effects of Removing All Subsidies
FAPRI (2002)	PE	-	+11.44% (removal of all distortions)
Goreux (2004b)	PE	-	+12% (removal of all subsidies) <i>Sensitivity Analysis: +2.9% to +13.4%</i>
ICAC (2002)	PE (ICAC Price & Demand Models)	+3c (5.7%) in 1999-00 +6c (10.5%) in 2000-01 +11c (26.3%) in 2001-02	+17c (29.7%) in 2000-01 +31c (74.2%) in 2001-02
ICAC (2003)	PE (ICAC Price & Demand Models)	-	+70% in 2001-02 and +15% in 2002-03
ICAC (2003a)	PE (ICAC Price & Demand Models)	-	+8c (14.4%) in 2002-03
Reeves et al. (2001)	CGE (GTAP) & PE (MFA Model)	Aust. Export Price +2.2% ¹⁵ (GTAP)	Australian Export Price +6% (MFA Model)
Tokarick (2003)	PE	-	+2.8% (removal of all support) +2.0% (removal of all production subsidies only)
Valderrama Becerra (2000)	PE (ICAC Price & Demand Models)	+3c (5.7%) in 1999-00 and +6c (10.5%) in 2000-01	-

¹⁴ Oxfam (2002) cites the estimates contained in ICAC (2002) for the removal of US subsidies only. Benin et al. (2003) appear to use the estimates from ICAC (2002) and a previous version of Goreux (2004b). Baffes (2004) cites a cross-section of the above estimates.

¹⁵ This estimate includes the impact of eliminating tariffs on raw cotton.

Table 4: Results of unit root tests on data series (1965-2001).¹⁶

Series	ADF (levels)	ADF (1 st diff.)	KPSS (levels)	KPSS (1 st diff.)
PW	-1.913626(t)	-6.387551***(t)	0.175912**(t)	0.066817(t)
	-2.274655(c)	-5.942562***(c)	0.367005*(c)	0.192335(c)
	-0.174648(n)	-6.013859***(n)		
WORLDPROD	-5.581868***(t)	-2.801339(t)	0.095783(t)	0.162448**(t)
	-0.509609(c)	-3.032923**(c)	0.705015**(c)	0.281820(c)
	2.092485 (n)	-7.509263***(n)		
WORLDCONS	-1.918261(t)	-5.759500***(t)	0.091524(t)	0.081313(t)
	-0.089473(c)	-5.792790***(c)	0.698963**(c)	0.096836(c)
	2.959305(n)	-4.807032***(n)		
DSTOCKS	-3.221192(t)	-	0.097770(t)	-
	-3.023897**(c)	-	0.186904(c)	-
	-2.624987**(n)	-		
TOTSUBLB	-3.827684**(t)	-7.031943***(t)	0.097105(t)	0.150803**(t)
	-3.856532*** (c)	-7.141310***(c)	0.175962(c)	0.154911(c)
	-1.433038(n)	-7.250313***(n)		

Table 5: Results of additional unit root tests on PW.

Series	ADF (levels, 1973-2001)	KPSS (levels, 1973-2001)	Perron (1965-2001)
PW	-3.604969**(t)		-3.395328(a) ¹⁷
	-3.528859**(c)	0.090592(t)	-2.687514(b)
	-0.244864(n)	0.205665(c)	-3.402809(c)

Table 6: Results of equation by equation diagnostic tests.

	LOG(PW)	LOG (WORLDPROD)	LOG (WORLDCONS)	DLOGSTOCKS	LOG (TOTSUBLB)
Equation S.E.	0.149298	0.047821	0.023750	0.136265	0.871019
R ²	0.925564	0.979625	0.994302	0.806060	0.760204
Adj. R ²	0.761804	0.934800	0.981765	0.379391	0.232652
Prob.(F)	0.003638	0.000009	0.000000	0.148963	0.280314
Jarque-Bera	2.756667	0.754717	1.130648	0.874450	0.312005
LM-1	1.156043	15.46085***	0.294526	26.32335***	5.343363**
LM-2	20.02442***	16.44996***	14.21083***	29.70841***	14.33488***
LM-4	26.32217***	26.16179***	14.89306***	30.11734***	26.68990***
Q-12	14.699	20.680*	9.2002	26.275**	7.6424
ARCH-1	5.609749**	1.004216	0.095200	0.013978	0.231288

Table 7: Results of system diagnostic tests.

Test	Result	Test	Result
Mod. (max char. root)	0.980737	Multivariate normal	120.5986
LM-1	26.42557	LM-2	48.62989***
LM-3	42.87478**	LM-4	44.53276**
LM-5	24.32657	LM-6	24.37422
LM-7	28.10533	LM-8	33.81688
LM-9	38.00767**	LM-10	32.29549
LM-11	50.35801***	LM-12	22.62838

¹⁶ Statistical significance is indicated with * (10%), ** (5%) and *** (1%). ADF and KPSS test specifications are indicated as follows: (t) includes both a constant and linear trend, (c) includes a constant only and (n) includes neither a linear trend nor a constant. Perron test specifications use the same letter codes as in Perron (1989): (a) allows for a broken trend, (b) for a mean shift and (c) for both. Lag lengths are selected so as to minimise the Schwartz Criterion (ADF and Perron tests) or using Newey-West automatic bandwidth selection (KPSS). For full test specifications see: Greene (2000), Kwiatkowski et al. (1992) and Perron (1989).

¹⁷ Models (a) and (b) only narrowly fail to reject the null hypothesis at the 10% level (critical values=-3.47 and -3.66 respectively).

Table 8: Forecast evaluation (static simulation, 1969-2001).

Test Statistic	Value	Test Statistic	Value
MAPE (PW)	5.763177	RMSE (PW)	5.006540
MAPE (PROD)	2.195128	RMSE (PROD)	2.050992
MAPE (CONS)	1.102175	RMSE (CONS)	0.961396
MAPE (DLOGSTOCKS)	66.73479	RMSE (DLOGSTOCKS)	0.075012
MAPE (TOTSUBLB)	48.42354	RMSE (TOTSUBLB)	7.509474

Table 9: Forecast evaluation for the world price (dynamic simulation, 1997-2001).

Test Statistic	Value	Test Statistic	Value
MAE	3.042921	MAPE	6.2416
RMSE	4.188650	RMSPE	0.950679

Table 10: Forecast evaluation (static simulation, 2002).

Test Statistic	Value	Test Statistic	Value
APE(PW)	15.6	APE(PROD)	5.533484
APE(CONS)	5.517605	APE(DLOGSTOCKS)	0.437383
APE(TOTSUBLB)	-		

Table 11: Results of block exogeneity (Granger causality) tests.

Equation	Exclude Price	Exclude Production	Exclude Consumption	Exclude Stock Movements	Exclude Subsidies	Exclude All
LOG(PW)	-	7.601576 ¹⁸	13.725363***	7.869595*	2.818494	19.373070
LOG (WORLD PROD)	34.44127***	-	31.581368***	35.775664***	39.317228***	55.733494***
LOG (WORLD CONS)	32.997320***	3.419394	-	1.314877	29.644686***	40.388336***
DLOG STOCKS	32.903292***	24.941755***	28.262833***	-	28.284378***	48.060954***
LOG (TOT SUBLB)	4.338281	6.239326	8.297239*	6.028363	-	20.379994

Table 12: Results of exclusion tests for US subsidies as an exogenous regressor.

Equation	LR Test
LOG(PW)	0.020075
LOG (WORLDPROD)	26.901558***
LOG (WORLDCONS)	15.0515669***
DLOGSTOCKS	34.108776***
ALL	9.705244**

¹⁸ In fact, the null hypothesis is very close to rejection at the 10% level, with a prob. value of 0.1073.

Table 13: Overview of Data and Sources¹⁹

Name	Coverage	Units	Description	Source
ELSPROD_t	1965-2002	K bales	US production of ELS cotton	USDA Cotton and Wool Yearbook, 2003, Table 3
ELSSUB_t	1968-2001	US\$K (curr.)	US government payments for ELS cotton	USDA ELS Cotton Factsheet, 1 January 2003, Table "Summary of Basic Data"
ELSSUBLB_t	1968-2001	US c/lb (curr.)	US government payments for ELS cotton, per pound of domestic ELS production	Calculated as (ELSSUB*10)/(ELSPROD*48)
TOTSUB_t	1964-2001	US\$M (curr.)	Total US government payments for ELS and upland cotton	Calculated as UPLSUB + ELSSUB/1000, setting ELSSUB to zero 1964-1967
TOTSUBLB_t	1964-2001	US c/lb (curr.)	Total US government payments for ELS and upland cotton, per pound of domestic cotton production	Calculated as (TOTSUB*10000)/(USPROD*48)
UPLSUB_t	1964-2001	US\$M (curr.)	US government payments for upland cotton	USDA Upland Cotton Factsheet, 1 January 2003, table "Summary of Basic Data" completed by "Government Payments by Crop Year"
UPLSUBLB_t	1964-2001	US c/lb (curr.)	US government payments for upland cotton, per pound of domestic upland production	Calculated as (UPLSUB*10000)/(UPLPROD*48)
UPLPROD_t	1965-2002	K bales	US production of upland cotton	USDA Cotton and Wool Yearbook, 2003, Table 2.
USPROD_t	1965-2002	K bales	Total US cotton production	Calculated as ELSPROD + UPLPROD and cross-checked against USDA Cotton and Wool Yearbook, 2003, Table 1
WORLDCONS_t	1965-2002	M bales	Total world cotton consumption	USDA Cotton and Wool Yearbook, 2003, Table 15
WORLDPROD_t	1965-2002	M bales	Total world cotton production	USDA Cotton and Wool Yearbook, 2003, Table 15
STOCKS_t	1965-2002	M bales	World beginning stocks	USDA Cotton and Wool Yearbook, 2003 Table 15
DSTOCKS_t	1965-2001	M bales	Movement in world stocks over current crop year	Calculated as STOCKS _{t+1} -STOCKS _t
PW_t	1963-2002	US c/lb (curr.)	Liverpool price	IMF International Financial Statistics, monthly data converted to crop year simple averages

¹⁹ All years refer to crop years, i.e. August to July. A bale is 480lb.

Figures

Figure 1: A basic partial equilibrium model of the cotton market with subsidies. (Source: Goreux, 2004b).




Figure 2: World Cotton Market - Quantities 1965-2002. (Source: USDA.)





Figure 3: World market cotton price (nominal and real), 1963-2002²⁰



²⁰ Real prices are deflated by the manufacturing import unit value, taken from Baffes (2004).

Figure 4: US cotton subsidies (nominal and real), 1964-2001.²¹





Figure 5: US cotton subsidies per unit of production (nominal and real), 1965-2001.



²¹ Real subsidies in Figures 4 and 5 are deflated using the US GDP deflator (source: IMF International Financial Statistics).

Figure 6: Cumulative IRF plots.

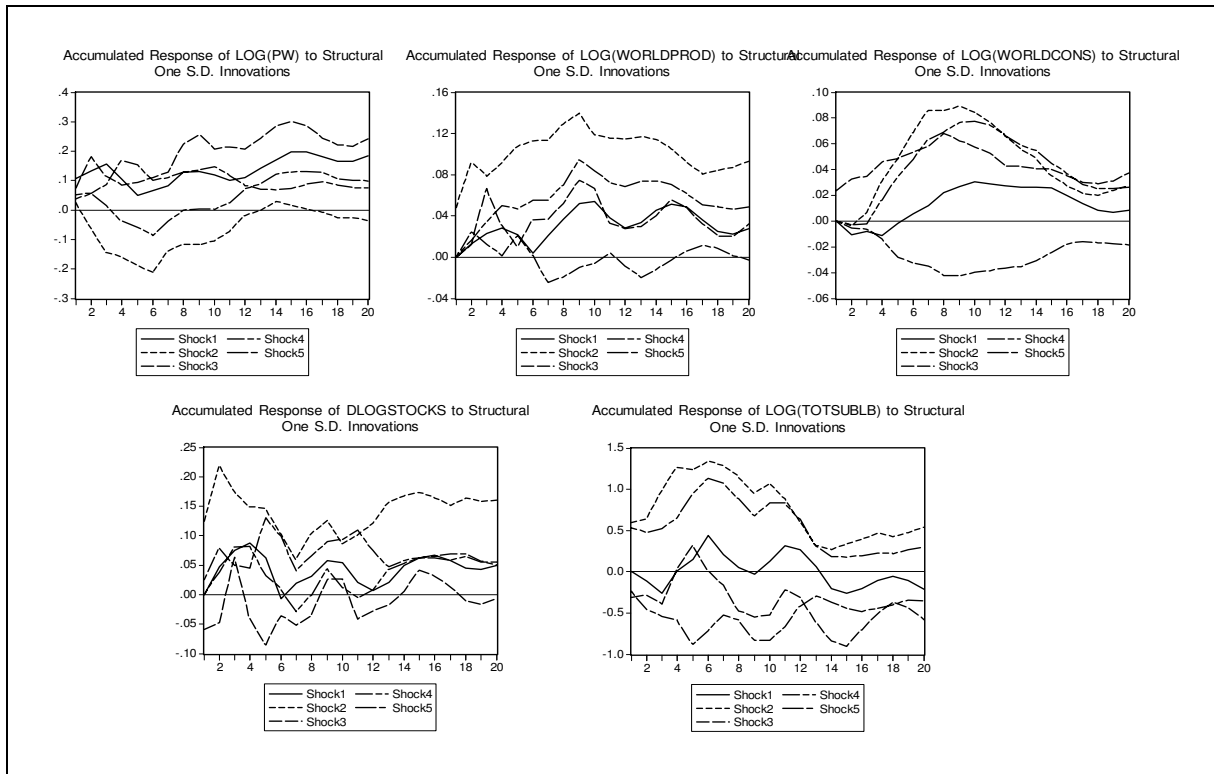


Figure 7: Cumulative IRF plots +/- 2 standard errors.

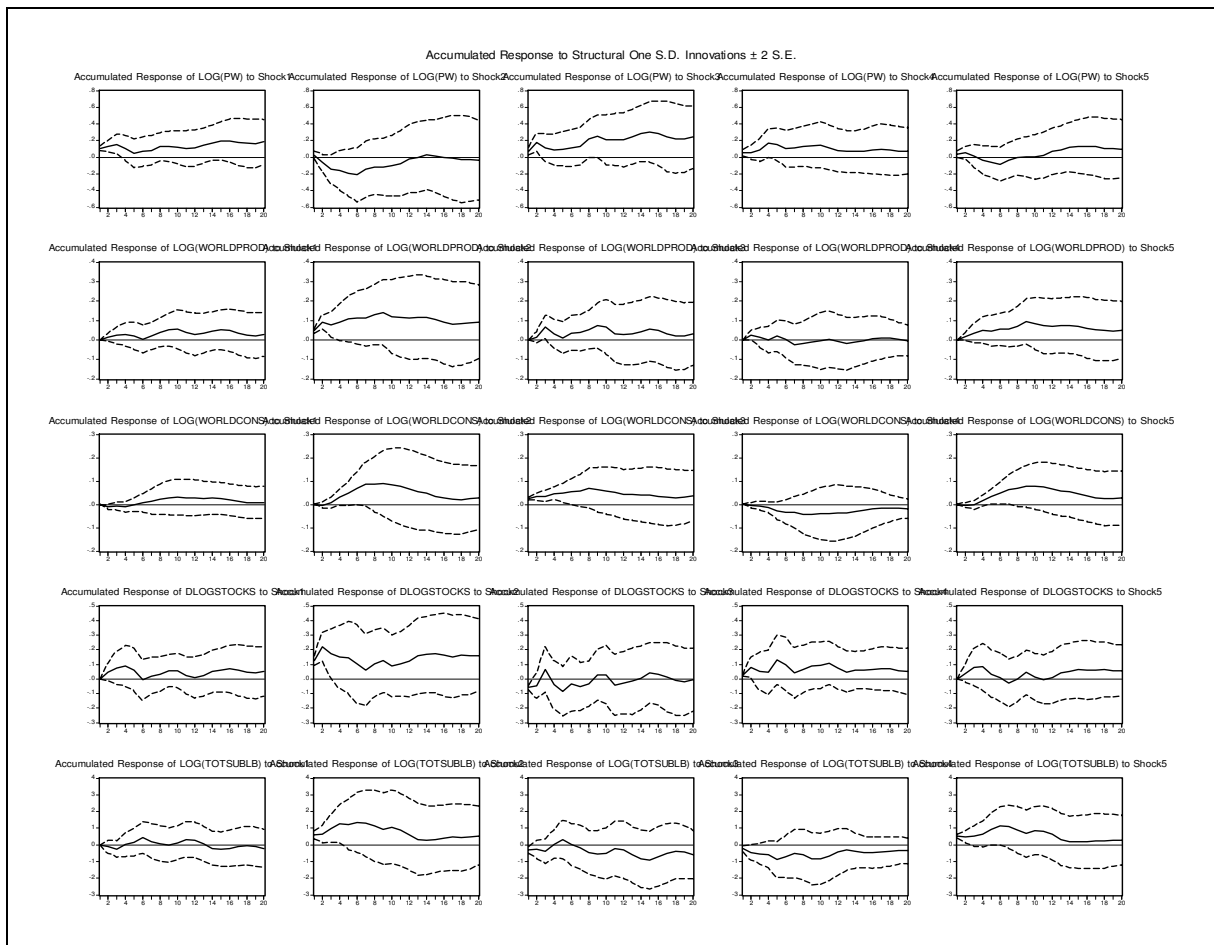


Figure 8: Forecast error variance decompositions.

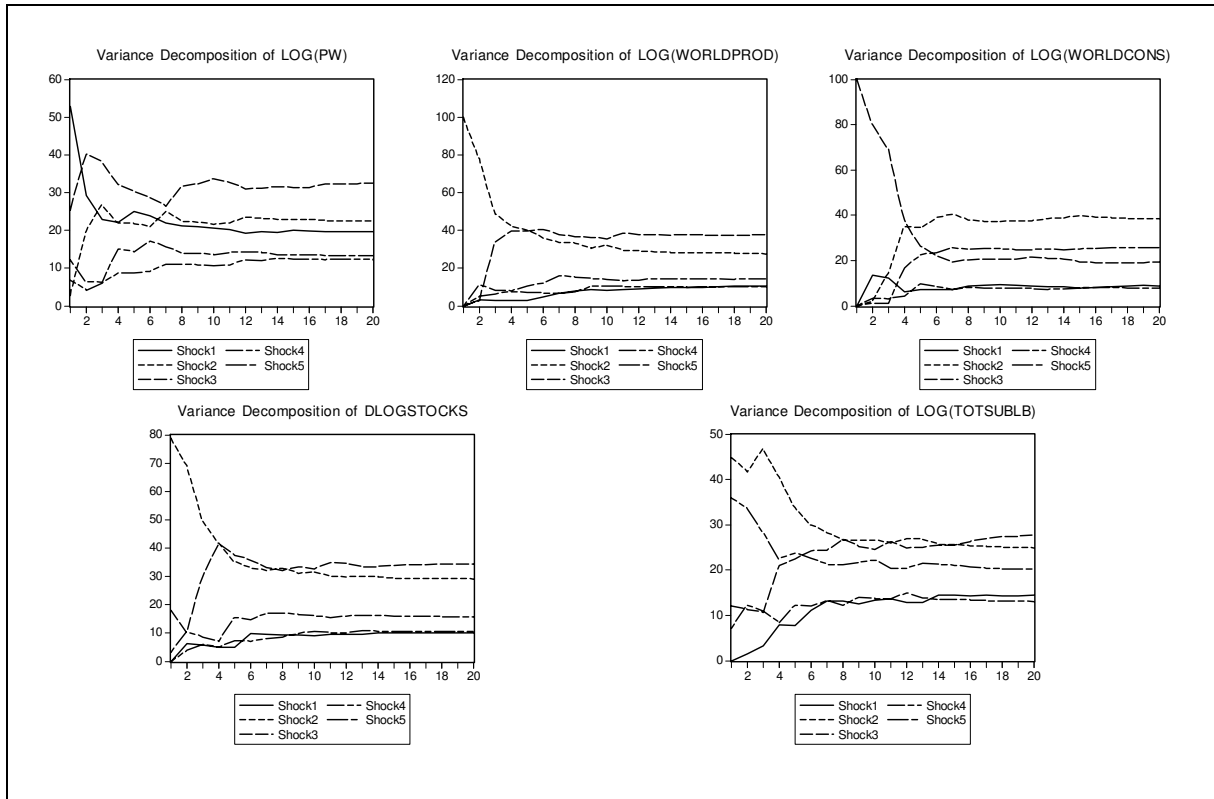


Figure 9: Simulation of a 10% reduction in US subsidies.




Figure 10: Simulation of a 50% reduction in US subsidies.

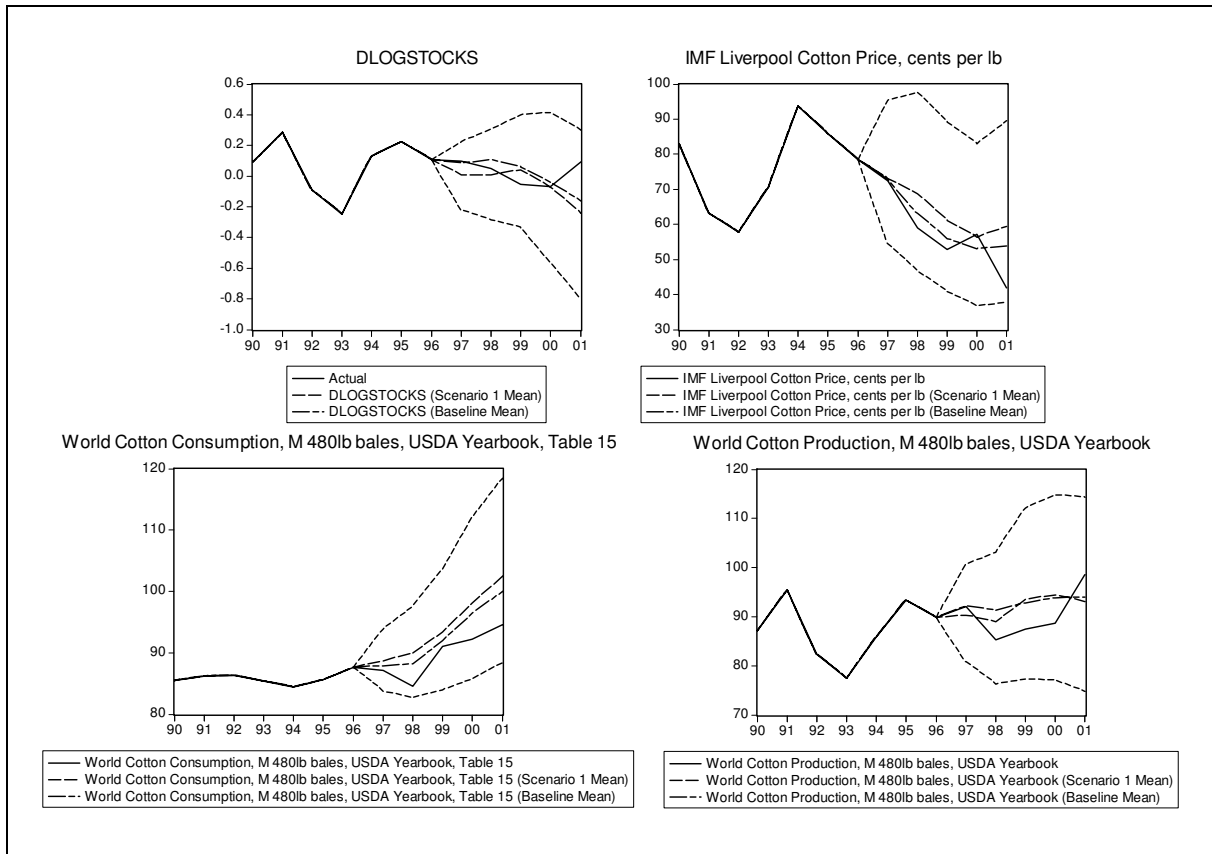


Figure 11: Simulation of a 90% reduction in US subsidies.

