AgEcon SEARCH
RESEARCH IN AGRICULTURAL & APPLIED ECONOMICS

The World's Largest Open Access Agricultural & Applied Economics Digital Library

This document is discoverable and free to researchers across the globe due to the work of AgEcon Search.

Help ensure our sustainability.

Give to AgEcon Search

AgEcon Search
<http://ageconsearch.umn.edu>
aesearch@umn.edu

*Papers downloaded from **AgEcon Search** may be used for non-commercial purposes and personal study only. No other use, including posting to another Internet site, is permitted without permission from the copyright owner (not AgEcon Search), or as allowed under the provisions of Fair Use, U.S. Copyright Act, Title 17 U.S.C.*

Discussion Paper 2015:4

Reducing unwanted consequences of aggregation in large-scale economic models - a systematic empirical evaluation with the GTAP model

Wolfgang Britz¹, Arne Drud² and Dominique van der Mensbrugghe³

1 Institute for Food and Resource Economics, University Bonn, Germany

2 ARKI Consulting & Development A/S, Bagsvaerd, Denmark

3 Center for Global Trade Analysis (GTAP), Purdue University, West Lafayette, IN, United States

The series "Agricultural and Resource Economics, Discussion Paper" contains preliminary manuscripts which are not (yet) published in professional journals, but have been subjected to an internal review. Comments and criticisms are welcome and should be sent to the author(s) directly. All citations need to be cleared with the corresponding author or the editor.

Editor: Thomas Heckelei

Institute for Food and Resource Economics

University of Bonn

Nußallee 21

53115 Bonn, Germany

Phone: +49-228-732332

Fax: +49-228-734693

E-mail: thomas.heckelei@ilr.uni-bonn.de

Reducing unwanted consequences of aggregation in large-scale economic models - a systematic empirical evaluation with the GTAP model

Wolfgang Britz, Arne Drud and Dominique van der Mensbrugghe

Abstract

We discuss how to avoid aggregation bias in large-scale global Computable General Equilibrium (CGE) models by reducing the need of pre-model aggregation, based on the combination of algorithmic improvements and a filtering approach which removes small transactions. Using large-scale sensitivity analysis, we show the impact of pre-aggregation and filtering on model size, model solution time and simulated welfare impacts, using a multi-lateral partial trade liberalization simulated with the standard GTAP model as the test case. We conclude that pre-model aggregation should be avoided as far as possible, and that our filtering approach and algorithmic improvements allow global CGE analysis even with highly disaggregated data sets at moderate solution times.

Keywords: Computable General Equilibrium analysis, aggregation bias

JEL classification: C68, C63

1 Background and Introduction

Despite considerable progress in soft- and hardware, computational issues such as solution time, memory requirements or numerical stability continue to constrain the use of economic simulation models. Therefore, the resolution of available data bases is typically not fully exploited in the application of economic models such as computable general equilibrium (CGE), supply side and partial equilibrium models. Rather, researchers aggregate data bases e.g. with regard to commodity and regional detail in order to yield models which can be solved in a fast and stable manner. The widely used GTAP data base (Narayanan et al. 2012), to give an example, is shipped with the GTAPAgg aggregation tool (Horridge 2006), and almost any

published application of the GTAP model (Hertel 1997) uses a study-specific pre-model aggregation of sectors and regions.¹ At the same time, it is well known that key model results such as welfare changes depend on the chosen aggregation level (cf. Ko and Britz 2013). However, almost no publication presenting and analyzing model results comprises a section with a sensitivity analysis with regard to alternative pre-model aggregations.

Aggregation issues have clearly gained interest in the community of economic modelers, also as a consequence of availability of more detailed data. A growing body of literature deals e.g. with the aggregation of trade policy instruments in economic modeling (e.g. Pelikan and Brockmeyer 2008) both with examples of combining different model types (Grant et al. 2007) or using welfare-consistent aggregators (cf. Himics and Britz 2015). So far, implementation of these approaches in larger modeling exercises is still scarce. And clearly, better aggregation of trade policy instruments can only improve one of the many possible sources of aggregation bias in modeling exercises. Our paper therefore adds a different perspective, beyond aggregation of policy wedges, by looking at impacts of pre-model aggregation. Additionally, it also presents and analyzes a complementary approach to pre-model aggregation based on filtering out small transactions while maintaining data consistency and important economic totals. We use large-scale sensitivity analysis to check how pre-model aggregation and our proposed filtering approach impact simulated results, but also model size and solution behavior. Besides complementing the discussion on appropriate aggregation in economic model, a further aim is to arrive at suggestions in applied CGE modeling helping practitioners in their daily work.

Our paper is structured as follows. The first section presents reasons for pre-aggregation, before we discuss possible approaches to reduce the need to do so. Our filtering algorithm is presented in section three, followed by an empirical evaluation. Finally, we summarize and conclude.

¹ This paper is focused on global economic databases and models where aggregation is most relevant. Single country databases and models typically have a resolution that can forego aggregation. Exceptions might be countries with highly detailed activities and commodities data bases such as the 500+ input-output table of the United States. Aggregation may also be useful for sub-accounts, for example a country database that contains a large household survey.

2 Reasons for pre-model aggregation of data bases and empirical evidence of its impacts

Pre-aggregation is a step before an actual application of the economic model which affects both its data base and parameterization. In applied trade modeling, pre-aggregation typically relates both to the commodity/sector and the regional dimension. Pre-aggregation can be understood as a multi-stage process. Most data are sampled originally at the level of individual agents and compiled by statistical offices to arrive at published totals. Typically, only these compiled totals are available to the research community; they suffer potentially from reporting and processing errors and might comprise systematic bias: observations might have been excluded due to data confidentiality issues or due to sampling thresholds, e.g. with regard to minimum firm size.

In order to make published totals suitable for modeling purposes, specialized teams next process them further, e.g. to render different data sources consistent to each other, to exclude outliers or to fill gaps. The efforts to release the GTAP data base provide many examples for such work such as the compilation of global bilateral trade flows and related data on trade protection levels. Most of the data sources used in that process need to be aggregated in order to arrive at the agreed upon regional and sector resolution of the final GTAP data base, covering 57 sectors and 134 regions in version 8 as used by us.²

The costs to duplicate such efforts in order to generate a differently detailed data set would be tremendous. The vast majority of economic modelers therefore rely on ready-to-use data sets such as the GTAP data base which are already pre-processed specifically for the purpose of economic modeling, but imply at the same time a specific maximal disaggregation of analysis. Generally, it would be desirable to exploit the full information of these data sets with regard to economic transactions, policy instruments and behavioral parameters, and increasingly also with regard to so-called satellites accounts with e.g. (bio)-physical information. Full information

² Issues of creating a globally consistent snapshot of the global economy go beyond simply reconciling official data from various sources (e.g. national statistical agencies, COMTRADE, IMF's Balance of Payments, etc.), but also reconciling national IO tables that are provided with different base years and in different currencies. All are adjusted to a common year and denominated in a common currency.

exploitation, however, requires running the model without any further aggregation which is often considered impossible for computational reasons – model size, solution time and stability – such that some form of pre-model aggregation is typically deemed unavoidable.

Besides reducing model size, which eases model solution and result analysis, pre-model aggregation has additional advantages. If errors in the data base are not or even negatively correlated across regions and sectors, aggregation will systematically improve data quality. Peaks in policy wedges as in cost and trade shares will be flattened which helps model solution. Aggregation will also systematically remove small entries such that the overall scaling of the model improves. It reduces the dimensionality of the simulated impacts which eases model analysis. If additional data are needed in a modeling exercise, e.g. environmental accounts, or different types of models are linked, pre-aggregation might also be necessary to arrive at common regional and product definitions.

Ideally, a selected pre-model aggregation should lead at the same time to an acceptable solution behavior and small aggregation bias, i.e. differences in simulated results compared to the model used in conjunction with the data base in full resolution. Unfortunately, as the full-fledged model can typically not be solved –at least not at acceptable cost–, the aggregation bias is not known. Accordingly, the researcher has to make an educated guess which aggregation by regions and sectors will cause a still acceptable bias, e.g. by combining sectors and regions with a similar cost structure or protection level. However, aggregation clearly offers quite limited degrees of freedom in that respect. To give an example: a set of regions belonging to one aggregate might have comparable protection rates in certain sectors and highly divergent ones in others. The flexibility in selecting the aggregation is further reduced if pre-model aggregation serves additional purposes such as to condense model results for easier analysis and presentation, e.g. by aggregating regions to continents or political blocks.

Clearly, study-specific, flexible pre-aggregation is only possible if the model is based on template equations which are structurally identical, i.e. do not vary in structure across regions and sectors, such that differences are solely depicted by parameterization. Accordingly, differently aggregated data sets change only the number of equations and the model's

parameterization, but not its structure. The remainder of the paper thus relates only to that rather common type of economic model. Basically all CGE models are based on equation templates, but also Multi-Commodity models such as IMPACT (Rosegrant et al. 2008), the COSIMO part of AGLINK-COSIMO (OECD 2007), CAPRI (Britz and Witzke 2014) or ESIM (Banse et al. 2005) from the field of agricultural economics. That does not imply that all these models could be easily run at a more aggregated or dis-aggregated regional and or commodity resolution of their current data bases. Aggregation e.g. of policy instruments or behavioral parameters might be a demanding task, and computer code to do so not available. Furthermore, the model's computer code might not be easily changed to work with different lists of products and regions.

There is ample evidence of unwanted impacts of using less detailed data bases in CGE modelling. Caron 2012 finds larger differences in estimated international carbon leakage effects based on same CGE model structure using a standard GTAP data base compared to one with more detail for industrial sectors. Alexeeva-Talebi et al. 2012 in a similar exercise focus more on the potential loss of detailed information when using more aggregated data bases, but also report aggregation bias at the aggregate level. Comparing two different sector aggregations and PE against GE closures, Brockmeier and Bektasoglu 2014 conclude that sectoral detail can matter substantially and turns out more important in their analysis than using a GE layout. Ko and Britz 2013 report larger changes in welfare gains from a EU-Korean free trade agreement when only the regional detail for the European Union is changed. All these findings suggest that at least certain types of pre-model aggregation can indeed systematically affect simulated results.

When looking at the reasons to pre-aggregate, only availability of additional data needed for a modeling exercise or model linkage might be a strict constraint, while computational barriers might be overcome.

3 Model setup and algorithmic improvements

Model solution behavior can clearly be improved by algorithmic changes. We discuss here some of the steps taken by us in that respect. But before, we briefly present our test framework.

In the following, we use a GAMS (Brooke et al. 2008) implementation of the well-known GTAP Standard Model (Hertel 1997). GAMS is a widely used language in economic numerical

simulation (cf. Britz and Kallrath 2012), and the implementation by van der Mensbrugghe and Britz (2015) used here offers those researchers who prefer GAMS over GEMPACK (Harrison and Pearson 1996) a fully compatible and tested version of the Standard GTAP Model.

The standard GTAP model and extensions thereof have been extensively used since the early 1990s for a wide range of different research and policy questions. The default in its application is a pre-model aggregation of the GTAP data base. Using the combination of this model and data base is therefore an obvious candidate for a systematic analysis of the impacts of pre-model aggregation. Furthermore, the standard GTAP model's basic features are quite common in global CGE analysis. On the production side, a Leontief technology for intermediate input demand and total value added demand is combined with a CES function substituting between primary factors under the assumption of constant returns to scale. On the demand side, substitution between domestic and imported goods in demand is modeled by a two-stage Armington structure. Markets are assumed to be competitive, and policies are expressed by relative price wedges. Different assumptions with regard to factor mobility are possible; we opt here for full factor mobility with the exemption of natural resources which are considered immobile. Further detail can be found in (Hertel 1997); a documentation of the individual equations as implemented in the GAMS version provide van der Mensbrugghe and Britz (2015).

We solve the model as a constrained system of non-linear equations (CNS), using CONOPT (Drud 1994) as the solver. A solution as a Mixed Complementarity Problem (MCP) is also supported by our GAMS code, but has shown to somewhat slow down model solution times. As we are not exploiting MCP features in our modeling exercise, we therefore only analyze the more standard CNS case where endogenous variables stay away from their bounds and equalities hold in the final solution. The model's equations are completely written in levels, and not as in the widely used GEMPACK version of the standard GTAP model as a mix of log-linearized and level equations.

In order to speed up solution time, we have firstly developed a sequential pre-solve algorithm which solves all single region models for several times at iteratively updated international prices before the full global setup. During these pre-solves, only one region is considered at a time with bi-lateral import prices fixed. Equally, the single country faces fixed aggregate import demand from their trading partners at fixed import prices of all other regions. Accordingly,

bilateral export demand for the region under consideration is reacting to changes in that country's fob prices due to updates of the import shares at the export destinations. Once all regional models had been processed several times in parallel using the grid solve algorithm from GAMS³ at updated international prices, the full model is solved. That procedure speeds up overall solution time dramatically for highly dis-aggregated model versions with many sectors and regions as it exploits knowledge about strong and weak relations between blocks of equations and variables in the overall global model which cannot be easily detected by the solver. It typically provides a very good starting point for the full model with very little infeasibilities left. For models with limited sectoral and regional detail, the pre-solves generate unnecessary overhead and can therefore be switched off.

Equally, we substitute out basically all linear definitional relations between variables in the model using the macro facility of GAMS⁴. That reduces especially model set-up time and memory use, but might have a limited impact on time needed by the solver depending on the solution algorithm used. CONOPT, to give an example, will detect definitional equations and treat them differently. Unfortunately, substitution of equations might affect solution behavior negatively, as the solver might introduce tiny infeasibilities at or below the feasibility tolerance in equations otherwise substituted out which relaxes the solution space. In our tests, a negative impact on solution time or stability from substitution could, however, not be observed. In order to reduce the impact of differently scaled economic transactions depicted by the elements of the aggregated SAMs and represented by variables in the model, we programmed code which defines scaling factors for all equations in the model, using the bigger of the value of the related transactions and a rather tiny threshold. That should ensure that smaller regions and sectors are

³ We use the beta release 24.6 of GAMS which allows for in-memory parallel execution of several model instances as the central part of its grid solve approach. Earlier version handled the instances via scratch directories which slows down parallel processing compared to the new approach especially in case that solution time for each instance has a small share in the overall handling of the model instances.

⁴ This substitution option likewise is available in the GEMPACK software (Harrison et al. 1994), which allows for the user to identify both the variables to be substituted out and the equation that can be used to as a substitute for the variable.

solved towards a similar relative accuracy as larger ones, and render the model better suited to work with more disaggregated data bases where differences in variable levels tend to be larger.

Another reason why researchers pre-aggregate is to ease results analysis (cf. Britz et al. 2015). Here, our code implementation covers additional flexible post-model aggregation (Britz and van der Mensbrugghe 2015) which aggregates over regions and sectors such that result analysis can be conducted either at the detail with which the actual model run is performed, or at a more aggregate, user chosen level.

Finally, we used a prototype implementation of CONOPT called CONOPT4 which has been updated to take advantage of the larger amount of memory on modern computers such that memory is no longer a practical limitation. In addition, CONOPT4 can use parallel execution. On demand, the inversion of the constraint matrix, the most expensive part of solving large CNS models, is to a large extent parallelized. Furthermore, the evaluation of the equations and their derivatives can be done in parallel using a new feature in GAMS. It should be mentioned that the parallel feature in CONOPT4 is only used for the large and expensive overall model (where it is most needed); the smaller regional models in the pre-solve algorithm use each one thread only as they are already run in parallel using the grid solve approach.

4 Filtering as a data driven alternative to pre-aggregation

4.1 Overview and motivation

An alternative or complement to an informed, research driven decision on a specific pre-model regional and sector aggregation is a data driven approach. We propose and present here an algorithm which aggregates small transactions in global SAMs to larger ones in a flexible manner. It combines three elements:

1. Filtering out small economic transactions from the data base
2. Rebalancing of the thinned out regional SAMs and bi-lateral trade flows
3. Maintaining closely the sum of important transactions such as GDP

That approach can be applied to global SAMs with different regional and sectoral detail and thus can be combined with pre-model aggregation as discussed above. A filtering algorithm

based on (1) and (2) has been distributed for some years to complement the GTAPinGAMS model (Rutherford and Harbor 2005), but we add here additional elements.⁵

Graph 1 below renders it obvious why it might pay off to remove small transactions: a large share of the about 2.6 million non-zero transactions found in the GTAP 8 data base is quite small. Changes in these transactions should not have a significant impact e.g. on simulated welfare impacts, but might affect solution behavior negatively, be it by simply driving up solution time and memory needs or even worse, by affecting solution stability.

Graph 1: Histogram with distribution of size of transactions in original SAMS / trade flows

Source: Own compilation

The data driven filtering approach will remove potentially a large share of these tiny values. Consequently, less aggressive pre-model aggregation is needed in order to arrive at models

⁵ The authors would like to thank Tom Rutherford not only for letting them use his filtering code as a starting point for the discussed extensions, but also for helpful feedback.

which can be solved in an acceptable time. That should systematically reduce the aggregation bias especially resulting from flattening out the distribution of policy wedges, as we will show in the empirical part of the paper. Equally, removing tiny values can reduce computational difficulties.

4.2 *Algorithmic details of the data driven aggregation process*

As indicated above, the algorithm consists of three major parts: (1) filtering, (2) re-balancing and (3) data driven aggregation, now discussed in some detail.

The filtering process tries to remove small items from the regional SAMs and from transactions relating to bi-lateral trade. The definition of what is small is based on a comparison relative to specific totals. In particular, domestic or imported intermediate demands of a commodity by a sector are measured against total intermediate demand of that sector. Demand by the government, for investment or by households for domestic produce or imports of a commodity are compared against the respective total domestic or import demand. Trade flows of a product are dropped if both shares on total exports of that product by the exporter and shares on imports of that product by the importer are below the relative threshold. Production is flagged as to be removed if net production of a commodity, i.e. after intermediate use of that commodity in its own production is deducted, is below the relative threshold with regard to total net production, or if both domestic and export demand have been flagged as deleted.

After introducing the flags which indicate data entries to remove, the regional SAMs are re-balanced, one region at the time. In order to ensure that bi-lateral imports and exports match, import and export trade flows have to be fixed during the re-balancing step of the individual regional SAMs. The balancing constraints for each regional SAM comprise three types of equalities. The first one ensures revenue exhaustion for each sector, i.e. the value of domestic and export sales net of output taxes of that sector must be equal to its production cost, i.e. the value of domestic and imported intermediates plus primary factor costs, taking into account any taxes. The second constraint refers to the domestic market: total domestic sales are equal to final, export and intermediate demand. The last of the balancing constraints is equally structured and relates to the import market. Additional constraints ensure that factor demands and total intermediate demand are not dropped if production remains non-zero.

Besides the fixed exports and imports, all other discussed elements are endogenous variables in the re-balancing problem which enter its objective function. Elements flagged as deleted receive penalty terms with a high weight to favor sparsity. The remaining elements are drawn towards the original data based on relative quadratic difference compared to the initial value. Using relative differences instead of absolute ones has shown to give similar results for larger transactions, but to improve the outcome for smaller ones: cost, consumption and tax shares tend to stay closer to the initial ones.

Without further modification, the discussed algorithm would have very little to do with aggregation issues per se. The process removes small data entries, such that the regional economies and international trade would systematically shrink. The algorithm is therefore expanded by two further elements. The first element relates to trade flows. Total trade is scaled to maintain the original world totals, and import and export tariffs are scaled to recover each country's original protection. That is achieved by first deriving a bi-lateral scaling factor which takes the change in total imports and exports for each country into account. It is applied to the individual trade flows. Afterwards, the corrected trade flows for each commodity are scaled to match the original world totals and tax rates on imports and exports are re-adjusted to match the original bi-lateral country totals of import and export tax revenues. Finally, transport margins are updated to recover the original transport demand per commodity and region. That process is applied before the actual SAM re-balancing, as trade related variables are fixed during the re-balancing step of the individual regional SAMs. The second element which contributes to maintain totals comprises constraints defining important totals such as GDP and total use of each primary factor. Deviations from these totals enter the objective function with a high weight.

These two additional elements imply a kind of data driven aggregation: the economic value of the dropped small transactions is not lost, but added to the remaining, larger transactions of the same type. The relative quadratic deviations in the objective function should distribute them rather equally in relative terms to these larger ones. The algorithm is implemented in GAMS and uses CONOPT as the solver.

The relative thresholds which determine the transactions to be dropped are introduced step-wise over iterations for two reasons. Firstly, as non-dropped transactions are scaled upwards, some

transaction might become larger than the thresholds over these iterations. And secondly, perhaps more importantly, it allows checking for a desired number of non-zero transactions and thus to control more easily the resulting model size.

4.3 *Exemplary results from filtering*

This section shows the impact of filtering on the GTAP data using a selected pre-aggregated data base. We opted to maintain the full sectoral resolution (i.e. 57 sectors) and tested different regional aggregations in the following. Generally, the largest reduction in the number of data entries is resulting from removing tiny intermediate demand transactions and bi-lateral trade flows.

Table 1 below gives an illustrative example from filtering a 57 sector x 56 region aggregation, i.e. the full sectoral detail of the GTAP data base is maintained. The unfiltered data based has around 800,000 non-zero transactions. It is first worthwhile to note that the regional SAMs and even the bilateral trade flows are almost completely dense in the starting point, i.e. after pre-aggregation of the GTAP data: the non-zero final demands are equal to number of sectors times the number of regions. There is an additional product (the trade margin) which enters the calculation, which explains why the amount of non-zero domestic output entries is somewhat higher. Also the intermediate demand matrices are completely dense. Even the trade matrix has with about 175,000 non-zero entries almost the full potential size of about 180,000.

Already a quite tiny threshold of 0.01% (itr2) relative to the respective total (see section 4.2. above for their definitions) removes about one third of all transactions, with the largest absolute reduction achieved by thinning out intermediate demand. That alone removes about 200,000 entries, another 40,000 entries are taken from the bi-lateral trade flows. Government demand entries are the most affected in relative terms. The final threshold of 0.1% used in our example removes about half of the transactions. As discussed below, the impact on simulated welfare changes is often quite small.

Table 1: Example for impact of filtering on the number of transaction

	start	itr2	itr4	itr6	itr8	done	delta	delta (%)
All transactions	816152	547377	479927	438547	405561	390097	-426055	-52.2
Domestic output	3248	3045	2947	2867	2785	2743	-505	-15.55
Factor demand	14390	13651	13302	13024	12724	12576	-1814	-12.61
Bilateral trade flows	175672	137990	121088	111482	103885	100013	-75659	-43.07
Trade margins	239802	221940	204039	189384	177150	171117	-68685	-28.64
Import demand sectors	185136	68134	52177	44233	38520	36139	-148997	-80.48
Domestic demand sectors	185136	96129	80586	72085	65246	62367	-122769	-66.31
Import demand Households	3192	2335	2177	2085	2023	1995	-1197	-37.5
Domestic demand Households	3192	2529	2335	2241	2160	2123	-1069	-33.49
Import demand Government	3192	958	753	676	627	602	-2590	-81.14
Domestic demand Government	3192	666	523	470	441	422	-2770	-86.78
Relative tolerance (%)		0.01	0.03	0.05	0.07	0.08		

Source: Own compilation. Note: uneven iterations removed to increase readability.

The user can run the filter program once with a rather aggressive threshold, say 5% and many steps (up to 30 are supported) to analyze the relation between the filtering threshold and the number of non-zero transaction left, but also to check which agents are affected most. The program also generates a table with details for individual sectors/commodities, not shown above. That allows for an informed decision which relative and/or absolute threshold might be deemed acceptable. The empirical part following next shows impacts on model solution and solving time both from pre-aggregating and filtering, but also on welfare impacts in for selected shocks. These findings can further help to decide on a suitable combination of a specific study.

5 Empirical evaluation

Our empirical evaluation is based on a typical counterfactual analysis with global CGEs, a partial multi-lateral trade liberalization, and runs it against a benchmark based on aggregating and filtering from the same global GTAP 8 data set. The analysis of the outcomes is split into two sections. First, we show how data-driven filtering impacts model size and solution behavior under different shocks, while the second section analyzes how pre-model aggregation and filtering impacts simulated welfare changes.

5.1 *The test framework*

We used for a test a two-step aggregation/filtering setup: first, the original data base is aggregated to a pre-defined level of regional detail using GTAPAgg, and next the filtering/aggregation algorithm is applied. Specifically, we generated data sets at the full 57 sector resolution with regional aggregates of 10, 16, 24, 36, 45, 57, 68 and 82 regions, using mostly population size as a guideline to select countries to single out in more dis-aggregated versions. We used filtering on these different aggregated versions of the GTAP data base to derive data sets between 60,000 and 400,000 transactions in steps of 20,000, plus four sets with around 500,000, 600,000, 700,000 and 800,000 transactions. Clearly, not all data sets are available in all resolutions: some data sets are already quite small to start with, and others would require very aggressive filtering thresholds to reduce the number of transactions below a certain size.

The actual active thresholds can be found in Graph 2. Each colored curve represents one level of pre-aggregation. As the relative filtering threshold is increased (the vertical axis), the size of the resulting dataset in thousands shrinks (the horizontal axis). To give an example: the purple line for the 57x45 case shows that a 0.1% filtering threshold ($\log_{10}$ is -1) reduces the size of the global data base from around 400,000 to about 280,000 non-empty transactions. Using 1% instead shrinks the data base to 60,000 non-zero transactions.

In order to generate our test data sets, we changed step-wise the desired maximal size of the non-zero transactions and let the program for each desired maximum iteratively increase the relative threshold until a data set was left with not more than the desired number of non-zeros. The relative filtering thresholds shown in the graphic below have hence been determined automatically by our algorithm. We reduced however the search time for the algorithm by first running a rough scan with a large relative threshold for each data set (an example is shown in table 1) to have an indication about maximum relative thresholds matching our different desired data base sizes in order to avoid too many iterations. Even if generating each data set requires maximally some minutes, repeating the exercise for our 90 examples made that preliminary check worthwhile. At the same time, restricting the spread of the relative search range leads to more fine grained changes in the number of non-zeros and thus to data bases sizes closer to our desired number of non-zeros.

Graph 2: Data base size in 1000 and related filtering thresholds in log10%

Source: Own compilation

Where possible, also an unfiltered version of the pre-aggregated data base was used, and additionally one with a very tiny filtering threshold of 1.E-6%. In total, we considered about 90 different variants of the same structural model layout which solely differ in the data bases, and almost all have the full sectoral detail of the GTAP data base.

The reader should note that the pre-aggregation process in GTAPAgg also automatically aggregates the behavioral parameters. Maintaining the full sector breakdown of the GTAP data base in most experiments has the advantage of limiting aggregation bias due to sector-specific parameter aggregation, as key parameters such as the substitution elasticities used on the CES production functions and the Armington nests are sector-specific, but identical across regions. For comparison, we also run tests with two exemptions: 10 and 18 sector resolutions combined with 68 regions respectively regional aggregates. These two additional data sets are not filtered.

Graph 3: Relation between model size (in 1000 equations), pre-aggregation and filtering

Source: Own compilation

Graph 3 above shows the model size - measured as the number of non-fixed variables in the model which is equal to the number of equations - as a function of the chosen level of pre-aggregation and filtering. The first interesting observation is that the number of variables in the model is generally smaller than the number of non-zero transactions in the data base: the variants with 160,000 non-zero transactions (horizontal axis) lead to models with about 100,000 variables (vertical axis). That can be explained by the fact, as mentioned above, that purely linear definitional equations such as those defining bi-lateral import prices, import and domestic prices for the agents and the Armington demand quantities for intermediates are substituted out of the model. Furthermore, transactions linked to taxing economic activities are typically substituted out from the model. To give an example: instead of having a variable which captures the value of taxing individual bi-lateral imports, that transaction is depicted in the model equations by multiplying the bi-lateral trade flow quantity variable with bi-lateral import price variable plus the fixed tax rate. These substitutions are obviously offsetting the effect from splitting the transaction from the data base into a quantity and price index plus introducing

additional variables, e.g. to define total intermediate inputs and total value added and related prices for each sector.

Graph 3 suggests a more or less linear relation between the number of SAM entries and model size, rather independent from the chosen pre-aggregation. For the larger regional aggregations, quite aggressive relative thresholds between 0.5% and 2% are necessary to yield smaller models such that a huge share of the original transactions needs to be deleted (see Graph 2).

To give an example: in order to solve the 57 x 82 case in less than 2 minutes for our test shock (see next section), it is necessary to generate a data set with less than 300,000 non-zero entries. As the unfiltered 57 x 82 case has close to 900,000 equations, that would require to remove a lot of detail from the data base and might lead to clearly unwanted structural changes e.g. with regard to important cost and trade shares. Such aggressive filtering is clearly not recommended for any serious model application, and as discussed in the next section, also not necessary.

5.2 *Model solution behavior depending on filtering and aggregation*

In order to test model behavior, we have selected a typical simulation exercise for global CGE models: a multi-lateral trade liberalization scenario which cuts all export subsidies and import taxes by 50%.

Graph 4 below shows the solution times in seconds⁶ for our multi-lateral trade shock. Overall, model solution times are quite modest even for rather large models. Not shown are results without using the pre-solves where far more time is needed for larger models; equally, solving the model as an MCP with PATH required some more time and is not discussed here. As seen,

⁶ The experiments were run independent from each other on a multi-core computing server with Xenon 3 GHz cores and fast hard drives. Times reported refer to cold start with GAMS, i.e. not using any restart files, and encompasses all steps until variables and equations were stored back to disk, including post-model processing. They always include a trial solve of the benchmark case to check that the model is correctly set-up. The same experiment with a 10x10 model without pre-solves takes less than two seconds on that machine. Run times on a modern laptop should be around twice as much, as long as the machine features enough memory and disk space.

the maximal time needed in our test was less than quarter of an hour, it relates to solving the largest model considered, the 57 x 82 case without any filtering. Actual solution for smaller models (up to around 200,000 entries) can actually be lowered further if the pre-solves are switched off.

A first interesting observation is the fact that once the filtered versions come close to the unfiltered ones, the almost linear relation between the number of transactions and model solution time is broken. That effect can be seen for the 57 x 36 case at around 340,000 transactions, and for the 57 x 45 variant even earlier around 320,000 transactions.

Graph 4: Solution time in seconds with differently detailed data bases and model configurations

Source. Own compilation

The most important observation is clearly that even very large models can be solved in our larger test shock in around 5 minutes. We also tested various other shocks (50% in consumption taxes, factor taxes, direct taxes; 20% increases in factor endowment and 10 % increases in total factor productivity) on a subset of our input data, and found roughly similar solution times as for the multi-lateral trade shock. We thus tend to conclude that there are no computational

reasons for high pre-model aggregations in analysis with the GTAP standard model or variants thereof⁷ as long as not a very large set of runs is performed. Clearly, if one performs large-scale sensitivity experiments e.g. with behavioral parameters, solution times for larger models might still be considered too long.

The interesting question here is naturally to what extent simulated results depend on the aggregation and potential filtering of the data base. Do we really gain a lot by using models which are harder to solve due to more dis-aggregated data bases which might potentially only add a lot of rather tiny values? We will shed some light on that question in the next section.

5.3 *Simulated welfare impacts*

The impact of the combination of pre-model aggregation and filtering on model size and solution behavior discussed so far is clearly of interest for modelers, but what matters most are differences in results. We compare here one sole number at continental and global scale which received quite some attention over time: the welfare gain from tariff liberalization. Our simulated gains are well in line with similar studies and subject to the same critic (e.g. Ackerman and Gallagher 2008). But we clearly do not aim at policy analysis or a discussion of potential weaknesses in our or similar simulation exercise. Rather, we use that central metric for a systematic look at aggregation bias.

A first illustrative example provides the results for an African total shown in graph 5 below, aggregated from differently detailed results for single African countries and country blocks depending on the aggregation level used. An obvious finding is that aggressive filtering has a detrimental effect on the simulated welfare effect. As seen for example from the 57 x 45 case, very aggressive and clearly not recommended filtering such as deleting 75% of the transactions

⁷ We also ran similar tests where we added simultaneously an implementation of GTAP-AGR and GTAP-AEZ.

reduces the simulated welfare gain by about the same relative amount⁸. But equally, reducing the number of regions in our analysis by pre-aggregation can almost half the welfare impact on Africa as a whole. The larger welfare gains simulated with more disaggregated regional data sets stem mostly from adding detail for Africa, and one might assume that any study with a focus on Africa would have used a database with similar regional detail.

Graph 5: Simulated welfare gain for Africa

Source: Own compilation

The differences between simulated welfare change for a world total are less striking, see graph 6 below, but the same pattern arises: more regional detail increases the welfare gains. The reader is however reminded here that the GTAP aggregation facility does not delete any trade and

⁸ One might assume that filtering mostly will remove tiny bilateral trade flows, but table 1 above suggests that the reduction is more evenly spread across the data base. The reader is also reminded that the filtering algorithm re-scales trade volumes and related tax income and maintains important totals such as GDP.

related taxation during aggregation. The reduction in simulated welfare hence stems from the fact that peaks in protection rates are flattened by averaging across regions, and not from converting former international trade into domestic sales.

Given these results, some patience seems to be recommended. As the impact of a chosen pre-model aggregation on the simulated results might be hard to predict, it seems wise to opt in doubt for a more disaggregated model, even if running it might take somewhat longer. However, moderate filtering in the range of 0.1-0.3% seems to barely impact the simulated results, but allows running models with far more regional and sector detail compared to unfiltered versions.

Graph 6: Simulated welfare gain at world level

Source: Own compilation

The most interesting finding seems to be that using very small filtering values up to 0.001% had no discernible impact on simulated welfare changes, but can speed up solution time

considerably. That type of filtering can hence be always recommended. We should note here that the user interface available with GTAP model in GAMS (Britz and van der Mensbrugghe 2015) allows setting specific thresholds for selected regions and sectors, which allows limiting or eliminating completely the filtering process for exemptions which might of special interest in the analysis.

6 Summary and conclusions

We presented a new algorithm useful for CGE modeling which filters out tiny and thus from an economic viewpoint hardly important entries from global data bases while maintaining consistency and important economic totals. Due to a sizeable reduction in the number of transactions, the resulting models are far smaller and consequently solve faster and are more stable compared to using the data base in its original unfiltered form. We test the algorithm with a fully compatible implementation of the GTAP standard model in GAMS which adds a pre-solve algorithm based on solving single region modules, helping to speed up considerably solution time for larger models and shocks. Equally, we use a new version of CONOPT which allows for parallel execution.

We show that the approach allows running the model with a far higher number of sectors and regions compared to typical GTAP applications published in the literature. In order to check impacts on modeling outcomes we first construct data sets using the default approach, i.e. aggregating the data base over sectors and regions without filtering any small values. Next, we use our algorithm to obtain filtered data bases of similar sizes which implies more regional or sectoral detail. For these about 90 data base variants, we run a partial multi-lateral trade liberalization scenario and compare simulated welfare gains, including those based on a rather large, filtered data set. This allows us to check how aggregation and filtering impact simulated results, i.e. to quantify the aggregation bias. Our results show the well-known effect that less detail reduces the overall welfare impact. As we could solve models with full sectoral resolution and up to 82 regions even without any filtering in less than 15 minutes, we conclude that pre-model aggregation by sectors can generally not be recommended and that larger pre-model aggregations by regions should be equally avoided.

Our analysis seems to indicate that, when using models of comparable size, the aggregation bias from the default pre-model aggregation is larger compared to using a more disaggregated data base with subsequent filtering leading to a similar size. That holds as long as the filtering thresholds are not too aggressive. We thus conclude that it is recommended to use our more data driven approach when building a tailored aggregated data base entering a CGE modeling exercises. Our tests seem to indicate that it is possible to run the GTAP model, even with additional modules, at full sector breakdown and high number of regions (up to 80) when small transactions are removed from the data base even under relatively large shocks.

The GAMS code of the filtering algorithm and to perform model runs along with a Graphical User Interface is available from the GTAP center.

7 References

- Alexeeva-Talebi, V., Böhringer, C., Löschel, A., & Voigt, S. (2012). The value-added of sectoral disaggregation: Implications on competitive consequences of climate change policies. *Energy Economics*, 34: 127-142.
- Ackerman, F., and Gallagher, K. P. (2008). The shrinking gains from global trade liberalization in computable general equilibrium models: A critical assessment. *International Journal of Political Economy*, 37(1), 50-77.
- Banse M., Grethe, H. and Nolte, S. (2005). European Simulation Model (ESIM) in the General Algebraic Modeling System (GAMS): Model Documentation. Berlin and Göttingen
- Britz W. and van der Mensbrugghe, D. (2015). "Interactive interface for an extended version of the standard GTAP model in GAMS," GTAP Technical Paper, forthcoming
- Britz, W. and Kallrath, J. (2012). Economic Simulation Models in Agricultural Economics: The Current and Possible Future Role of Algebraic Modeling Languages in: Kallrath, J. (eds.): *Algebraic Modelling Systems: Modeling and Solving Real World Optimization Problems*, p. 199-212, Springer, Heidelberg, Germany
- Britz, W. and Witzke, P. (2014). CAPRI model documentation 2014, University Bonn, Institute for Food and Resource Economics. http://www.capri-model.org/docs/capri_documentation.pdf (last accessed 16.09.2015)

- Britz, W., Pèrez Dominguez, I., Narayanan, G. B. (2015). Analyzing Results from Agricultural Large-scale Economic Simulation Models: Recent Progress and the Way Ahead. *German Journal of Agricultural Economics* 64(2): 107 – 119
- Brockmeier, M., & Bektasoglu, B. (2014). Model structure or data aggregation level: Which leads to greater bias of results? *Economic Modelling* 38: 238-245.
- Brooke, A., Kendrick, D. and Meeraus, A. (1988). GAMS: A User's Guide. The Scientific Press: California.
- Caron, J. (2012). Estimating carbon leakage and the efficiency of border adjustments in general equilibrium—Does sectoral aggregation matter?. *Energy Economics*, 34: 111-126
- Drud, A. S. (1994). CONOPT—a large-scale GRG code. *ORSA Journal on Computing*, 6(2): 207-216.
- Grant, J. H., Hertel, T. W., and Rutherford, T. F. (2007). Tariff line analysis of US and international dairy protection. *Agricultural Economics* 37(1): 271-280.
- Harrison, W. Jill, and Ken R. Pearson (1996). Computing solutions for large general equilibrium models using GEMPACK. *Computational Economics* 9(2): 83-127.
- Hertel, Thomas.W., Ed., (1997). Global Trade Analysis: Modeling and Applications. Cambridge University Press, Massachutes, USA
- Himics M. and Britz W. (2015). Flexible and welfare-consistent tariff aggregation over exporter regions. *Economic Modelling*, first revision in review
- Harrison, W. J., Pearson, K. R., Powell, A. A and Small, E. J. (1994). Solving applied general equilibrium models represented as a mixture of linearized and levels equations. *Computational Economics*, 7(3), 203-223.
- Horridge, M. (2006). GTAPAgg data aggregation program. Global Trade, Assistance, and Production: The GTAP 6 Data Base.
- Ko, J.-H. and Britz, W. (2013) Does Regional and Sectoral Aggregation Matter? Sensitivity Analysis in the Context of an EU-South Korea FTA". <https://www.gtap.agecon.purdue.edu/resources/download/6313.pdf> (last accessed 16.09.2015)

- Narayanan, G.B., Aguiar, A. and McDougall, R. (eds.) (2012). Global Trade, Assistance, and Production: The GTAP 8 Data Base, Center for Global Trade Analysis, Purdue University
- OECD (2007). Documentation of the Aglink-Cosimo Model. Organization for Economic Co-operation and Development: Paris.
[http://www.oecd.org/officialdocuments/publicdisplaydocumentpdf/?cote=AGR/CA/APM\(2006\)16/FINAL&docLanguage=En..](http://www.oecd.org/officialdocuments/publicdisplaydocumentpdf/?cote=AGR/CA/APM(2006)16/FINAL&docLanguage=En..) (last accessed 16.09.2015)
- Pelikan, J., and Brockmeier, M. (2008). Methods to aggregate import tariffs and their impacts on modeling results. *Journal of Economic Integration* 23(3): 685-708.
- Rosegrant, M. W., Msangi, S., Ringler, C., Sulser, T. B., Zhu, T., and Cline, S. A. (2008). International model for policy analysis of agricultural commodities and trade (IMPACT): model description (p. 42). Washington, DC, USA: International Food Policy Research Institute.
- Rutherford T. and Harbor A. (2005). GTAP6inGAMS: The Dataset and Static Model. Prepared for the Workshop: "Applied General Equilibrium Modeling for Trade Policy Analysis in Russia and the CIS", The World Bank Resident Mission, Moscow, December 1-9, 2005.
<http://www.mpsge.org/gtap6/gtap6gams.pdf> (last accessed 16.09.2015)
- Van der Mensbrugghe D. and Britz W. (2015). The standard GTAP model in GAMS Version 6.2, GTAP Technical Paper, forthcoming