Ideas for the New NEMS Liquid Fuel Market Model

Frederic Murphy
Fox School of Business Temple University
Philadelphia PA 19122
fmurphy@temple.edu

Prior to making modeling decisions on the LFMM, EIA has requested suggestions for the model redesign. The views expressed in this paper are entirely mine and are not an official position of EIA. The intention is to kick off discussion, not finalize a design. To make the recommendations presented here, I looked at the documentation of current models, drew on my experiences with the previous EIA models, and visited the stakeholders in EIA.

Current Problems and Concerns with the Refinery Model

Based on EIA's discussion in the white paper, a key component of the Petroleum Market Model (PMM)[1], namely the refinery model, has been unsatisfactory for the past 3 to 5 years. Part of the problem is that the model is asked to answer almost unanswerable questions. The most recent experience with product and crude prices is illustrative of the problems. From the peak in 2008, crude prices fell to less than 1/3 the peak price in early 2009. In the same period, gasoline prices fell by little more than half, radically changing the spread between crude and product prices. At the peak, diesel and home-heating-oil prices were way above gasoline prices. A few years ago refinery margins were strong (large). Then they narrowed dramatically. The crude market was in a position of contango[2] in early 2009 to where it made sense to store crude in tankers in US costal areas and take arbitrage profits.

Refinery models essentially provide the spreads between crude and product prices. The current refinery model is a linear program. These spreads are quite volatile and there is a real question as to whether any forecast several years out is meaningful. The prices for crude are more stable on a percentage basis than the spreads for several reasons. First, there is a co-product pricing problem with petroleum products that does not exist with other forms of energy and the co-product prices are highly sensitive to small shifts in the demand mix and product imports. Second, product markets are truly world markets and a purely domestic model is driven by assumptions on product imports and exports, the world mix of crude qualities and exchange rates, not by interactions with other sectors of NEMS. The increased demand for distillate in China prior to the Olympics drove the spreads and crude prices in 2008, pushing diesel prices above gasoline. A question that has to be answered from the currently available data is whether there is the information to support reasonably accurate modeling of world supply, crude types, demand, and refinery capacities and capabilities.

A counter to what I have written is that in a longer-run model like NEMS, the spreads among products are capped by conversion units and those caps should provide a reasonable guess for the

[1] PMM is the predecessor to the proposed new Liquid Fuel Market Model (LFMM) in NEMS.
[2] Contago it is the situation where the price of a commodity for future delivery is *higher* than the current spot price, or a far future delivery price is higher than a nearer future delivery of the product.

spreads. Still, given that spreads are much more volatile than the underlying prices, estimating refinery profitability with any accuracy will be nearly impossible.

Another problem with the model is the regulations mandating minimum quantities of biofuels. As those minimums bind, they distort the prices and understate the total bill for petroleum products. To achieve full cost recovery, the duals in the current LP cannot be used as prices with fixed-quantity lower bounds. (Minimum percentages in a blending constraint will work.) The reason for this is that the duals on the lower bounds absorb the higher costs of the biofuels and those costs are not passed on to the duals that represent fuel prices.

The history of modeling the liquids market and the issues with NEMS convergence point to what not to do. The original PIES model used a full-scale refinery model to generate extreme points of input and output crude and product mixes. The extreme points were then inserted into the PIES LP. The problem with that approach is that there was no generation of added extreme points as needed as in the classic decomposition algorithm. The result was product prices that were so volatile as to be unusable as the fuel mixes moved away from the guessed extreme points. Generally, the product prices were buried in reports and not highlighted.

The next attempt was more successful at producing stable product prices. A base input-output vector was constructed from historical patterns to establish baseline quantities. Then a set of shift vectors where each vector moved (shifted) one product into another, while retaining a Btu equivalence, were constructed. The objective-function coefficient on a shift vector set the maximum price difference between the two products in the vector and fixed the price difference when that vector was basic.

This model made the prices behave. However, the model needed recalibrating every time demand changed significantly and there was no way to tie the PIES solutions to what was happening in refineries. That is, nothing affecting refinery policy could be modeled in PIES, making this approach inappropriate for NEMS. Still, it was the best model of all that had been tried for many years.

The virtue of the current NEMS refinery model is that one can look at refinery-specific issues. However, the model is large relative to the information it provides and it likely has many hidden bounds inserted to make the results look good. The model is just too hard to use without a lot of expert modeler intervention. Furthermore, the lack of full coordination with the world petroleum model means that tuning the model as it is whipsawed by imports is hard. Because the prices in a co-product pricing problem are so dependent on the mix of inputs and outputs, reasonable prices are not possible without capturing all of the relevant interactions. Half measures on estimating the mix of crudes and the kinds of product imports will lead to overly rigid quantities and unstable prices.

There has not been a complete rewrite of the model since the early 1990’s and it has passed through many different modeler’s hands. It currently consists of 45,000 lines of FORTRAN. Since 90% of that code is in the revise, it is clear that the current model is a layering of fixes to the original model. The revise should be used only for model elements that are scenario dependent, not the core structure of the model. With so many lines of code, the model is opaque; it violates the spirit of EIA’s mandate to be the open, transparent forecaster; and is error prone simply because of its size. The model has gone against modern software engineering.

An immediate problem is that the model is not written in a modern LP modeling language (matrix generation language)[3] that makes the model structure and data transparent. The reliance on Ketron software is risky because there has been no significant investment by Ketron in the software in a very long time and the key people in this part of the company are well beyond retirement age. Essentially, this aspect of the business is not understood by the owners of Ketron and they are not making the proper investments.

In the current and future resource-constrained environment of EIA, whatever model is built has to be simpler and easier to use with staff time devoted more to analyzing results and data and understanding markets, rather than wrestling with code. This means that the process for operating the model has to be simplified.

The structure of NEMS and its predecessor IFFS was chosen to simplify the process of making a forecast relative to PIES. Some of the design imperatives and computing technologies have changed over the past three decades, making it possible to improve the overall design of NEMS and improve the processes for building and operating the NEMS components. Furthermore, the new approaches make model documentation simpler. The refinery model stands out as an example of what can be changed. Modern database and modeling technologies should be part of any discussion of the new model and the processing steps and people hours required to generate a forecast should be an explicit part of the design of the new model, given the pattern of overwork at EIA.

A short summary of what the model should produce and how it should operate.

In making the list of issues and questions the model must address, I looked at the EIA assessment of the areas that should be addressed. I picked out the most desired and combined some in the list. The topics can be organized into a hierarchy from the narrowest to broadest questions as follows.

- Industry—understand the future state of the industry, especially in a low/no-growth future.
- Energy markets—forecast prices and quantities of petroleum-based fuels in the US.
- Energy policy/regulations—estimate the impacts of proposed policies and regulations on energy markets, including new sources of liquid fuels.
- Environmental policies and other non-energy policies—estimate the impacts of these policies on energy markets.

Although not something that the model should produce, an important feature of the new model should be that it is usable and easily understood by knowledgeable people both within and outside of EIA.

[3] Every modern LP package comes with a matrix generation language; see GAMS for an example one. During the initial development of the PMM, it was decided by the original developers that it would be more efficient to generate the model matrix outside of NEMS, read it in once at the beginning of a NEMS run, and then make "minor" and year-dependent revisions as the market and integrated solution changed. The revisions to the original matrix now overwhelm the matrix formulation portion.

In this paper I focus on the following forecast elements that are high in EIA's ranking.

- Refinery margins.
- Liquid fuel production, consumption, and imports.
- Impacts of alternative fuels on production, consumption, imports and refineries.
- Impacts of tax and tax-subsidy policies.
- Refinery investments and the future of the industry.

Refinery margins are listed first because producing realistic margins is a problem with the current model.

Knowability

The first administrator of EIA, Lincoln Moses, used to say, "There are no facts about the future." Given the data quality in that era, we also used to say "there are no facts about the past either." Central to building a new model, we need to understand what can and cannot be known with any reasonable precision.

We explore this in the context of refinery margins, which are an important issue for EIA. What should be knowable from the model, given imports, the available crude slate, the refineries' capital stock, and demands is the expected margins. Furthermore, varying imports, crude slates, and capacities should lead to shifts in margins in the model that make sense a priori or add insight through an understanding of the model behavior. Thus, the model should backcast margins reasonably well. If the model backcasts well and margins behave appropriately with sensitivity runs, the forecasted margins can be considered reasonable estimates of expected margins under normal operating conditions in the future. However, given that chaotic events are always happening, EIA needs to establish broad bands on the margins.

What the model cannot do is capture the actual variability in margins. This depends too much on events beyond the resolution of NEMS, such as refinery fires, hurricanes, etc. that would skew the prices above the normal-market estimates. Thus, the model will likely underestimate refinery profits coming from normal random events in the industry that usually take out capacity and raise prices. Also, EIA will miss transient, unprofitable periods. The August 2009 market is illustrative. Heavy crudes were deeply discounted in 2008 as they were the only source of increased production. In 2009 the spreads have narrowed greatly because, as demand has fallen, heavy crude production has fallen more than the lighter crudes. Currently, heavy crude units have marginal economics on an operating basis.

The import slate depends on the differences between production and demand elsewhere. The domestic refinery model cannot provide intelligent margins without a carefully built international refining sector, world crude supply by crude type, and international estimates of petroleum-product demand. Furthermore, product/blend stock movements contain noise based on the sequence of transactions involving those products and the physical product moves along the supply chain. The EIA data on sources of products show product coming from places that should not supply the US on a regular basis. Understanding this data would help to understand what the model is saying. Although there are anomalies in the data, the patterns generally follow least cost.

Given the ability of an LP basis to shift significantly from a cost change of a few cents, it is not clear that the model can correctly capture the effects of the world on the US portion of the refining industry with any stability in the quantity flows. The least-cost flow patterns are subject to not only refining but also shipping costs that change the relative economics of moving crude and product. The Baltic Dry Index crashed to 10% of its peak in 6 months during 2008. Tanker rates followed a similar pattern. A VLCC daily rate in early 2009 fell to around what a 50,000 barrel barge was getting in early to mid 2008. These rates determine the international movements of crude and product as much as refinery-unit economics. This is especially true with Aframax tankers in the north Atlantic trade between the US and Europe that carry RBOB and crude.

Sometimes it is hard to choose what should go into the design of scenarios. Usually it is some external factor that throws off forecasts. These need to be understood for developing the range of uncertainty. Here is an example of a factor that is not addressed by any of the EIA models. Although crude prices are denominated on dollars, crude is a world commodity. The volatility of the dollar has affected import prices for crude and products relative to the rest of the world. The current models are not designed to deal with exchange rates. However, they will be important in determining consumption in the rest of the world versus the US. With the dollar having to fall relative to a basket of Asian currencies, the extent of the currency movements will play a major role in pricing crude in the US. Product prices are also affected as the non-US refining costs depend on exchange rates.

Some of these problems should be dealt with through scenario design. EIA cannot be in the business of forecasting tanker rates but can examine a range of transportation costs. A scenario with the dollar devalued so that the balance of payments reverts to 0 would lead to high US prices and lower rest-of-world prices. By keying the scenario to balanced trade, EIA can remain "policy neutral" but capture a real phenomenon that is central to the forecasts.

Modeling choices

The modeling choices that can be made fall into two broad categories, modeling the decisions with the intent of capturing the consequences and modeling the consequences of the decisions directly. Process models represent a stylized view of the world within which the decisions must be made and then the decisions are made within the model. The most common approach here is optimization modeling as a simulation of a market, with the optimizer functioning as a purely rational decision maker. When modeling the consequences of a process or industry, one gathers data on the industry outputs and inputs, posits a set of relationships between the outputs and inputs as equations, and estimates a set of parameters that link the inputs to the outputs.

The dichotomy is not firm in that activities in linear programs are black boxes of inputs and outputs and the more detailed the structural equations in an econometric model the closer the representation of the decision. When there is no data history, as in the case with new technologies, then modeling the process is the only choice. However, when there is a data history, a process model should be able to reproduce history.

LP process models have the problem that they can change their solutions radically with minor perturbations in cost. The standard fixes are to grow the model until it is unwieldy or put in "judgment" bounds that make the solution look good but restrict the model to the point where it fails to capture actual flexibility in the industry modeled. Growing the model gives a false sense of precision because the odds of an input error increase rapidly with the increased size and the elements added do not necessarily have a good data foundation. This is a problem with the current refinery model.

An alternative to an LP is to use sharing functions. I personally dislike using sharing functions except in very limited circumstances. The reason is that the sharing functions too often have their parameters set without using data. This is the same as setting bounds in an LP to give a good answer. However, the impacts of the parameter choices are harder to discern because the model results are not overly rigid.

Econometric models have their weaknesses too. The data series are limited, leading to problems with the significance of the estimations and ability to represent enough of the structure to capture what is really going on. They cannot capture readily policies involving new technologies or regulations that have no data history. Lastly, because the decisions are more implicit than explicit, it is harder to justify model results.

The two approaches have virtues and faults. The art of modeling involves being thoughtful on choosing between the two approaches, deciding how the two approaches can be blended in a larger model, understanding the balance between detail and usable information, and designing a model that is operational in the sense that it can be exercised enough so that its behavior is fully understood and a new policy or scenario can be implemented without undo pain.

Ideas

Positioning the model within NEMS.

The LFMM includes crude and NGL production, imports, exports, the refinery model, and petroleum product consumption. Consumption takes place in the end-use sectors, residential, commercial, industrial, and transportation. The demand elasticities of the end-use sectors in any given year are small and the cross elasticities are even smaller. The two sectors where fuel switching within a solution year is potentially important are in electricity generation and refining. Petroleum-product use in electricity generation represents 1.5% of petroleum consumption and is confined mainly to special circumstances where fuel switching should not be a big factor. The refinery model captures the tradeoff between using oil and natural gas for refinery fuel and this is also not a large number. Furthermore, the switching between oil and gas can be controlled based on an equation that estimates the shares based on price differences between oil and gas, data that is readily available within EIA.

The interactions with the other sectors in NEMS take place over time as capital-stock turnover and the accumulation cross elasticity effects with other fuel prices change the end-use fuel mix over time. The real interactions are with the rest of the world. This argues for more attention to world markets in LFMM and less effort in the NEMS linkages. For example, the refinery model probably needs to run only once or twice in a solution year.

If the quantities shift more than I expect within a solution year, then a mini-model can be used in the larger NEMS to facilitate convergence only (and not used for any public forecast). The mini-model would be calibrated to be very close to the solutions from the larger model as part of the NEMS iterations. One possibility is to use the shift-vector approach described above with the base quantities and the relative product prices coming from runs of the refinery model. Another is to construct extreme points to represent the LP dual rather than the primal and iterate enough in a decomposition approach to get close using the extreme points.

The main point here is that interactions of domestic product markets with the rest of the world are far greater, and more important, than with the other domestic sectors in NEMS. Thus, the design effort should be on a balanced, integrated view of world petroleum markets, not complex linkages with the rest of NEMS.

LFMM choices

The key first decision is that the LFMM model, whatever it is, has to be one model for all aspects of forecasting and policy analysis activities of EIA to minimize the cost of coordinating two models covering the same sector. Otherwise, too much staff time is spent fiddling and the smaller model tends to drive the results, leading to too many imposed bounds on the larger model.

- Annual operations of refineries, US. Use a linear programming model of refineries.
    - Pros
        - It's what refiners do, so this is a simulation of what is actually done.
        - New technologies can be implemented directly and the behavior of the technologies can be observed directly.
        - Duals make good prices, except when regulations mess up the prices, as with biofuels.
        - Econometric modeling can model the relationship between product mix and price. However, this approach cannot handle biofuels or other regulatory changes.
        - Nonlinear models better fit some operating units. However, that adds too much complexity.
    - Cons
        - It is too easy to make the model too large.
        - May get too much bang-bang in the solution from small cost changes, leading to administrative bounds or excessive growth in size.
    - What to do to make the decision
        - Run extensive experiments with the current model to see how well it backcasts and determine why or why not it does well.
        - Test to see what the contribution is of each structure in the model to the prices and quantities to determine if the added complexity has value. The goal is to see how big or small each black box should be, remembering that complexity is not free.
            - Included in the testing should be the current structure of two refineries per region, which seems to be redundant, and the different processing units, which seems to be excessive detail.
            - Refinery complexity by region should be varied to see where it adds value.
        - See what data is available for model building and benchmarking.

        - At what detail are capacities and unit operating behaviors available?
        - What kind of benchmarking data is available in terms of historical time series?
    - § Search the research literature on aggregate refinery modeling for best practices. I quickly found "Financial Risk Management with Product Pricing in the Planning of Refinery Operations" by Lakkhanawat and Bagajewicz in *Ind. Eng. Chem. Res.* **2008,** *47,* 6622–6639 by using scholar.google.com. Using the cites feature and other search techniques of this website it is possible to find out what the literature is.
    - § Determine the industry practice on aggregate modeling of US petroleum markets, as the purpose of the model is to simulate industry behavior.
    - § Make a modeling language decision
        - Modern modeling languages make for more transparent modeling when there are good implementation practices. They also eliminate a large amount of labor in the mechanics of doing a run.
        - The language should make the model transparent.
        - A model with 45,000 lines of FORTRAN is too hard to run and understand and too error prone.
        - EIA faces serious supplier risk. Everyone central to OML is over retirement age and there seems to be no investment in new development or commitment to the product beyond the legacy staff. EIA will have to replace OML at some point. It is better to do it on EIA terms now rather than as a crisis later.
        - Since it is too hard to do a model run with the current system, there is probably not enough experimenting to get a real feel for how the model behaves.
- o Miscellaneous
    - § Use the current PADD structure as the data that EIA collects matches that breakdown (I think). Since PADD 3 effectively drives the relative prices of products in almost all regions east of the Rockies through basic flow activities, more attention to detail in that region over others would provide a good modeling balance.

Domestic capacity expansion. The issue of foresight in capacity expansion has been a problem for years. Everyone agrees that perfect foresight is wrong. At the same time, adaptive foresight produces wiggles in the numbers and staff time is devoted to dampening the wiggles. Furthermore, there is no theoretical justification for the approach. With the current financial-engineering failures, we are finding that we even do not have reliable probability distributions. Managers have known this for years. For capacity decisions, it is better to model the outcomes of decision processes rather than the decisions. When there is no relevant data, myopic decisions can be better than artificially constructed foresight because at least it is clear how the decisions are made in the model. Businesses use payback time and high hurdle rates to compensate for lack of reliable knowledge of the future, which is not done in NEMS. The current refinery model uses too much heavy crude. This means that either there is no modeling of the capacity for processing heavy crude or capacity can expand either myopically or too cheaply to process all of the available heavy crude.

- There are two possibilities for modeling capacity expansion and retirements: embedding capacity activities in the refinery LP and estimating capacity expansion econometrically.

    - Why important capacity activities should not be embedded in the operating LP.
        - Expensive units are not added in reality by using one run of a deterministic LP.
        - The bang-bang properties of the solution lead to "administrative" bounds that are too inflexible.
        - The problem of foresight and rational expectations in a deterministic model are well known.
        - Even though the quantities are quite stable within a year, doing full NEMS iterations to address foresight leads to price forecasts that would oscillate too much and the cost of doing this is greater than the value.
        - Modeling the expansion of inexpensive units in the LP, if their capacities are modeled at all, is acceptable because they have quick payback and there is no issue of foresight.
    - Why capacity activities are embedded in the LP.
        - It's easy in the beginning.
        - It keeps the modeling methodology the same.
    - Why capacity expansion should be modeled econometrically.
        - Even though refiners use LP's for these decisions, they run many sensitivities and make internal assumptions about future crude prices and product slates. This means LP is a computational tool but not the decision maker. Different companies have different forecasts of future commodity prices, demand levels and crude slates. This results in a distribution of decisions, not all-or-nothing outcomes. An estimated model captures the effect of the distribution of outcomes, unlike capacity activities in an LP.
        - Even though units such as heavy crude units are discrete lumps of capacity when added, at the scale of EIA's forecasts that is not a big deal and half a unit is effectively an expected value.
        - Capacity additions would be smooth and respond appropriately to crude price differentials.
        - Capacities are only indirectly linked to demands because of product imports. So, the data series used should include net demands on US refineries.
    - What to do to test the ideas.
        - See what data is available.
        - See what the literature is.
        - Understand industry practice.
        - Try some estimations and see what the fits look like. If they don't look good, embed the capacity decisions in the LP. When using the LP model for capacity planning, admit that the expansion plans cannot be realistic unless fixed in advance, which is a non-data-driven estimation.
        - I am not clear as what the right independent variables should be. For heavy crude units, I would begin with overall refinery profitability and the price difference between heavy and light crudes in the current and several preceding years. Refinery capacity has grown at refineries that have not

closed down through debottlenecking. So, there is a time trend component. The crack spread in current and past years should be important as well.

- International refining operations. What should be estimated from the data and what should be modeled using an LP?
  - Pros for expanding the domestic linear program
    - The product market is a world market with some regional differences and the interactions of the domestic market with the world market are greater than the interactions with other markets in NEMS.
    - Conceptually, this is easy.
    - International refiners run LP's.
    - For the major international refining regions there probably is sufficient data to build LP's.
    - The international product slate and demand patterns complement the US demand patterns and product slate, leading to what should be a relatively stable, but evolving, world product mix.
  - Cons
    - If the models in each region are too small, then the imports patterns could be subject to wide swings.
    - If the models are too large, then LFMM becomes too unwieldy.
    - There may be data issues in developing an appropriate representation.
  - Pros for estimated import equations
    - The model stays small.
    - No specific policies need to be represented. So, modeling the decisions is less important.
  - Cons
    - If embedded directly in the refinery LP, the model becomes a nonlinear program, which is harder to solve. The alternative of iterating adds another step to NEMS, increasing run times.
    - The data may not support estimating what is happening in the BRIC nations (Brazil, Russia, India, and China)..
    - There are important policies in other countries that affect the US, for example, dieselization in Europe, that are not captured in the data as a policy consequence.
    - The model may not capture that gasoline is becoming a byproduct of refining in important regions of the world and the US has been the dumping ground for that gasoline.
  - What to do
    - Run tests with the current model to evaluate the balance between the domestic and international components of the current model and see if the model can backcast the international flows and prices.
    - Evaluate what components of the current international portion of the refinery model add value and where the representation is too approximate and needs further work.
    - Evaluate what data is available.
    - Examine the literature to see what the state of the research literature is.
    - Talk to people in planning at international oil companies to see what they do.

        - See if it is possible to do a rudimentary econometric fit to replace the LP submodels. If no clear patterns emerge from the data, quickly move away from further estimation attempts.
        - Make sure any model is not so rigid that the domestic portion of the refinery model has to absorb all of the adjustments in the shifts in world crude supply and product demand.

- International refinery capacities
    - Keep this very simple but do sensitivities to be sure that this part of the model does not drive the rest of the LP solution.
    - Try not letting the capacities be determined as part of the international operations submodel because that might give too much flexibility to add capacity internationally and dump product in the US.
    - There are two difficult components to address from a US perspective. The first is refinery closings in Europe and their effect on Northeast gasoline supply as dieselization continues. The second is the expanding capacity, especially by the BRIC's and what their product import needs or export behaviors will be. These are data question in the early years. Discerning what they will do in the long run is hard to say. The sensitivity of the model to a range of their behaviors should be explored.

- Crude supply. One of the main reasons why the model performs badly in terms of crude selection and co-product pricing is that there does not seem to have been a concerted effort to study of the mix of crudes available on the world market and how that slate will shift as production waxes and wanes in various regions. The relative product prices are as dependent on the crude slates as the unit capacities and no meaningful forecasts of margins and relative product prices can be done without including a good representation of the crude slate.
    - Look at the data on crude qualities coming from different petroleum provinces, both in the US and around the world. Combine that with information on the future profile of production in these provinces to come up with a reasonable future slate.
    - The different qualities used for the international crudes versus domestic crudes either double the activities for distillation yields in the operations model or the model does not represent appropriately the yields. Examine the value of having a different slate of crude qualities in the US versus the rest of the world and what should be the number of crudes using the operations LP as a test bed. (This is the kind of task that is hard to do now but is easy in a modern modeling system such as GAMS, MPL, or AMPL.)
    - During the recent market highs for crude, Saudi Heavy was deeply discounted. In August 2009 Saudi Arabia reduced production of heavy crude to stay within its OPEC quota and maximize revenues. This means that all crudes are not equal in constructing a supply curve and world spare capacity is most likely in heavier crude. This needs to be explored because I suspect that the mix of crudes is in fixed proportion on each step of the supply curve.
    - One of the recent market anomalies was the lower price for WTI than Brent. A factor in this, from what I understand, is the lack of processing capacity for crude from tar sands in the Midwest because of air-quality restrictions and the need to move that crude to the Gulf in the pipeline system. As part of any sensitivity

analysis this must be addressed in NEMS if the environmental concerns lead to permanent flow patterns with crude flows crossing while moving in opposite directions.

- Petroleum product demand

  - Domestic demand forecasts in the end-use sectors seem to be working reasonably well. However, there are some critical coordination issues that need to be addressed. I was told that the hybrid-vehicle model does not allow E-85 fuel. This means that as hybrids gain market share in the out years, the regular fleet is forced to take an outsized portion of the ethanol requirements. This imbalances the model. If what I was told is true, the assumptions of the models needs to be documented together in one place and a report writer written that highlights the details of the fuel use to eliminate this kind of unintended consequence of assumptions
  - One effect of the evolving non-US consumption is that the US is becoming the dumping ground for excess gasoline production. Do a study to understand how the product mix demanded in the various countries will shift. Is gasoline going to be a byproduct rather than the central product, as it is becoming in Europe? Or will the gasoline share rise in countries like India and China as autos become more important and kerosene for lighting and cooking becomes a smaller share of the barrel? I don't think that simple mix extrapolations will work here. Included in the analysis is whether the electricity supplies are getting more reliable or will diesel generators operate at greater levels in developing countries. This is important in determining the slate of products exported to the US.

- New technologies and market-altering regulations. There are four dimensions to this, subsidies and taxes, mandates, technology representation, and carbon policy.
  - Subsidies and taxes can be combined, as they represent adjustments to the cost coefficients on the activities. Care has to be taken in tracking them.
    - In the case of ethanol, the supply cost needs to be reduced by the subsidy to get the economic tradeoff correct in blending. The blending should take place in the demand regions with the appropriate transportation costs, by the way. Since the credit is applied to supply, the full gasoline tax should be applied to total production to track total costs properly and give the right marginals. The same rules apply to all taxes and subsidies.
  - Mandates can distort prices.
    - If a mandate is a minimum share of the product volume, then the duals on the outputs of the blending activities fully capture the increased cost of the mandated blending component, as the dual on this constraint is a measure of the increased marginal cost of the blend requirement for an increase in total fuel. The same applies to a minimum blending constraint where the minimum is a given proportion of demand.
    - If a mandate requires a minimum total quantity as a policy target, then the cost of the biofuel is not captured properly by the duals in the model when setting prices. The reason is that the dual on the lower bound absorbs the excess cost of the biofuels and these costs are not transmitted by the LP to the final product prices. To adjust the prices in this case, the model needs

to add a markup on the prices of the products containing that blend component equal to the dual on the minimum-quantity constraint times the resulting fraction of the mandated biofuels. This matches the cost "absorbed" by the lower bound with the incremental revenue required to cover those costs.

- Technology representation in the main is no different from other process technologies. The only potential difference is the representation of learning. However, this is no different from other new technologies elsewhere in NEMS and can be dealt with by adjusting costs using a learning curve only as the model moves forward in time.
- Carbon taxes and cap and trade cut across all fuels and apply to refinery fuel use. Carbon taxes are just markups. Cap and trade requires implementing a carbon constraint across all fuels. Having a refinery model that does not connect tightly into NEMS may seem problematical. However, it is not a problem. Since the demand elasticities are so small, the LFMM model can be solved with an initial guess at the carbon costs. Then the carbon content can be passed to the NEMS carbon constraint as a fixed quantity. The end-use prices should be adjusted by the difference between the initial guess and the carbon price for that year. With a sharing function for refinery fuel use, the quantity of carbon emitted would be relatively stable, and any volume changes should not be not a large number.
- What to do
  - From a theoretical or NEMS-structure perspective there is little new here. Consequently, what most needs to be understood is why the technologies seem to be behaving badly in the model.

- Other issues

  - The international transportation costs for product should be higher than crude to reflect the capital and operating costs of different vessel sizes. They should be a function of steel prices and fuel costs. The reason for including steel is that an oligopoly is forming around iron ore and EIA should look at the relationship between steel and energy prices, if there is one. If Commerce develops steel-price projections, include them in constructing transportation rates and other capital costs.
  - Essentially, NEMS covers a piece of the resource puzzle and the current forces that lead to high prices for energy will increase the prices of other commodities used in the energy sector as we face a more resource-constrained world. I used to think that high oil prices would dampen world economic activity, leading the prices of other commodities to move in the opposite direction of energy, as world oil depletes and the world economy becomes less resource intensive. This is definitely not true during business cycles like the one we are in because shifts in aggregate demand affect the equilibrium prices of all commodities in the same direction. However, given the strong economic growth taking place outside the developed world, the overall growth component on the commodity demand equation will mean that in equilibrium, the depletable commodities like certain minerals and non-depletable, but land-using commodities like agricultural products will have prices that correlate with oil, with the underlying cause being economic growth outside the developed world. What this implies for the macro

models and oil demand is unclear. Nevertheless, the US economy will become less resource intensive for non-energy commodities and that should also reduce the energy intensity even beyond what is being seen now.

- Exchange rates are an important driver in energy costs and are now very unstable. Scenario design in establishing crude prices and international refinery operating costs has to take this into account.
- In sum, energy has been treated within NEMS as a unique commodity that has to be understood because of the importance of energy to the US economy and was special because of outside non-market forces. The latest run up in world prices was the first since WWII not associated with a war in the Middle East but part of a general commodities boom. The meaning of this is unclear, except that energy is somewhat less special than other commodities and EIA's analyses need to take this into account.

Conclusions on the LFMM design

My view of what to do does not involve radical change from the current modeling approach. In terms of the form of the model, the main points are, make LFMM a world LP model, remove the important capacity and refinery fuel-consumption decisions from the LP to stabilize those numbers and better fit the data, and adjust the prices by the dual of any minimum biofuels requirement if necessary. Furthermore, the representation of the qualities of crudes has to be improved because there is no way that a refinery model can construct reasonable prices without both the right demand mix and the right mix of crudes.

That I recommend using linear programming to represent daily operations does not mean the model should not be rebuilt. I suspect that the main problems with the refinery model, outside of crude quality, come from using OML and not rebuilding the basic LP in years. Layers of fixes have probably taken over the model structure and are probably the source of instability and non-intuitive results. Because Ketron is a fading company and other modeling languages are more transparent, I strongly recommend moving from OML to a new modeling language. When moving to a new language, the structure of the database (table structure) needs to be worked out so that the model is comprehensible and multiple scenarios can be run easily, unlike older versions of Markal. The LFMM is a good candidate for beginning the transition from OML, as it is the LP that is least tied in to the iterative structure of NEMS.

The current model is a black box because of its layers of revisions. It is also a black box because it is too hard to exercise to better understand how it behaves. Whatever is done with the new model, it needs to be backcast and flexed thoroughly so that users have confidence in it. It should be amenable to quick experiments with regular backcasting to test the model.

Before fixing on a detailed implementation, an operations workflow diagram should be written of the current system. This includes all of the steps necessary to get a refinery run done in NEMS. By doing this, the time wasters in generating a run will come clear and they can be removed in the redesign. The main reason for moving from the PIES to the IFFS modeling structure was that PIES had an extremely inefficient workflow structure that burned through people. I used workflow diagrams to understand that redesign. That the same people have been able to work on IFFS and NEMS for years shows the structural redesign achieved it goal of providing quality of work life. Workflow diagrams would help achieve new efficiencies in

NEMS that would make jobs easier and increase the time allocated to the more interesting activities beyond getting out another run.

Further Thoughts

Several issues raised here relate to NEMS in general. First, the role of the systems view in model design should be revisited. PIES was developed as one large model because of a perfect storm hitting all energy sectors simultaneously. After the legislation of the 1970's was passed, energy policy moved to sector-specific issues and away from a broad systems view. An integrated forecast was still necessary because of the interactions among the different energy markets. IFFS and NEMS were designed around that less integrative view of policy analysis with separate sectors tied together in an iterative process to produce a unified forecast. That separation ended with the SO_2 program and the work needed to get NEMS to converge. Even though medium-term energy legislation is generally sector-specific, the environmental programs are not. These programs are examples of the way in which non-energy programs are more important than energy programs in shaping energy markets. Mercury and CO_2 require tightly integrated analysis tools of the energy sectors that interact strongly, electricity, coal, and natural gas. NEMS needs some design changes that are big from a conceptual perspective but relatively easy to do from an algorithmic perspective, using current models. There have been substantial advances in decomposition in equilibrium models that can be incorporated into NEMS that would speed convergence, make the model simpler to operate, and provide better and more stable integrated analyses of programs that cut across sectors such as carbon policy.

That energy markets are impacted more by non-energy policies means that EIA should have a role in capturing the effects of those policies on energy markets. The highway program had more impact on energy consumption than any energy policy. With the raft of new programs coming through other agencies, EIA should play a greater role in understanding their impacts, especially in the very long run. EIA should seek funding to expand its analysis capabilities in that direction, leading to more modeling and analysis beyond NEMS.

A real effort should be made to institutionalize a process for assessing how well the model fits history without administratively inserting bounds. One of the EIA people I interviewed said that in a prior job, a group he worked in was able to backcast its refinery model. However, once calibrated from the backcasting, the forecasts still missed the mark. This says there are important issues in defining the base case. This is a useful insight. However, it does not mean that EIA should throw up its hands and stop forecasting. It means that more has to be done to understand forecast uncertainty. Rather than just rely on sensitivity cases to establish uncertainty bounds. EIA should backcast its refinery model using assumptions that match what happened in that period and get the model to match what happened. Then the model should be rerun for those years using base case assumptions. The solution differences in each year represent the effects of the base-case assumptions and underlying randomness that cannot be worked out of the system. These differences should generate a distribution that can be used to establish uncertainty bounds beyond the standard ones used in scenario design and inform the scenario design process. The problem with this kind of randomness is that it can be explained in hindsight so that the forecast looks to be in error. However, these events are rarely forecast in advance. Thus, hindsight presents a false sense of determinism, as described in the book *The Black Swan*.