An Overview of the Biomass Scenario Model

July 2006 – July 2011

Steve Peterson
Lexidyne, LLC
Colorado Springs, Colorado

NREL Technical Monitor: Brian Bush

An Overview of the Biomass Scenario Model

July 2006 – July 2011

Steve Peterson
Lexidyne, LLC
Colorado Springs, Colorado

NREL Technical Monitor: Brian Bush
Prepared under Subcontract No. LGJ-1-11948-01

NOTICE

This report was prepared as an account of work sponsored by an agency of the United States government. Neither the United States government nor any agency thereof, nor any of their employees, makes any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States government or any agency thereof. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States government or any agency thereof.

This report is available at no cost from the National Renewable Energy Laboratory (NREL) at www.nrel.gov/publications.

Available electronically at SciTech Connect http:/www.osti.gov/scitech

Available for a processing fee to U.S. Department of Energy and its contractors, in paper, from:

U.S. Department of Energy
Office of Scientific and Technical Information
P.O. Box 62
Oak Ridge, TN 37831-0062
OSTI http://www.osti.gov
Phone: 865.576.8401
Fax: 865.576.5728
Email: reports@osti.gov

Available for sale to the public, in paper, from:

U.S. Department of Commerce
National Technical Information Service
5301 Shawnee Road
Alexandria, VA 22312
NTIS http://www.ntis.gov
Phone: 800.553.6847 or 703.605.6000
Fax: 703.605.6900
Email: orders@ntis.gov

Cover Photos by Dennis Schroeder: (left to right) NREL 26173, NREL 18302, NREL 19758, NREL 29642, NREL 19795.

NREL prints on paper that contains recycled content.

Acknowledgments

Over the course of this project, the project team has developed numerous internal reports and briefing documents. These are housed on the project repository at http://www.nrel.gov/analysis/bsm/. This report draws from these reports and briefing documents—particularly for the graphics used in this document. Versions of these graphics and a parallel description of much of the model also appears in (Newes, Inman, and Bush 2011), written by members of the BSM team.

Contents

Figures and Tables

1 Introduction

This report describes the structure of the October 2012 version of the Biomass Scenario Model (BSM) in considerable detail, oriented towards readers with a background or interest in the underlying modeling structures. Readers seeking a less-detailed summary of the BSM may refer to Peterson (2013). BSM aims to provide a framework for exploring the potential contribution of biofuel technologies to the transportation energy supply for the United States over the next several decades. The model has evolved significantly from the prototype developed as part of the Role of Biomass in America's Energy Future (RBAEF) project. BSM represents the supply chain surrounding conversion pathways for multiple fuel products, including ethanol, butanol, and infrastructure-compatible biofuels such as diesel, jet fuel, and gasoline.

1.1 Purpose and Strategy

Throughout the BSM effort, the BSM team–including staff from the National Renewable Energy Laboratory (NREL), United States Department of Energy (DOE) sponsors, other contractors, and Peterson—has worked to develop a framework for understanding the evolution of the supply chain for biofuels. The model offers a framework for understanding the circumstances that govern the rapid introduction of biofuels into the United States transportation fuel system. Figure 1, shown below, depicts our strategy and approach:
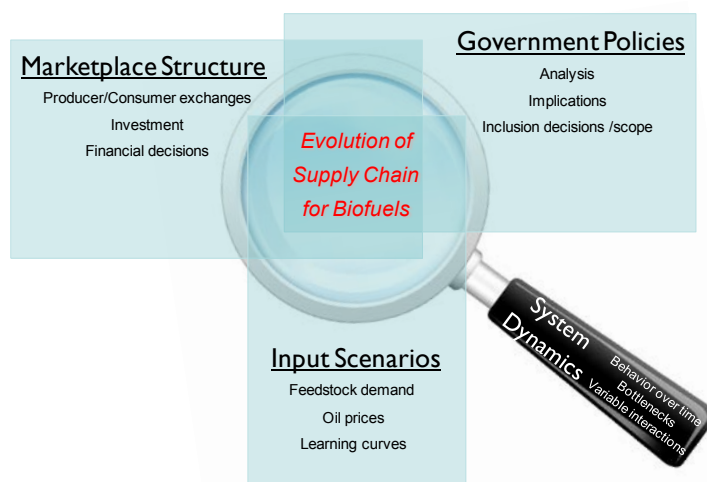
Figure 1. BSM strategy and approach

In order to gain a clear view into the evolution of the supply chain for biofuels, BSM focuses on the interplay between marketplace structures, various input scenarios, and government policy sets. In order to represent this rich interplay, we have employed system dynamics. System dynamics has a long history of application in a broad array of application areas (Sterman 2000). Its strengths in representing and simulating the behavior of feedback-rich social systems make it well-suited to the task of emphasizing and understanding system interconnections within and among the different stages of the biofuels supply chain.

1.2 What's in this Document

This report provides a description of the structure of BSM. It begins with an overview of the model architecture. Second, it takes a deep dive into the sectors and modules that comprise the

1

model. Moving across the supply chain, this second section provides a view into a) feedstock supply and logistics; b) multiple conversion pathways; c) "downstream" ethanol and butanol dynamics; d) the vehicle scenario module; e) the BSM representation of the oil industry; and f) ethanol import. Third, the report identifies key connections *between* model sectors. Fourth, the report briefly describes our approach to data inputs for the model. Fifth, a brief section describes directions for analysis with this version of the model. Appendices to the report provide detail on the use of logit formulations in BSM, on form of the pricing structure used within multiple components of the model, and on the mathematical approach used to create a simplified vehicle vintaging structure.

Three special types of diagrams are contained in this report:

- Traditional influence diagrams, such as seen in Figure 5, in which calculations/outputs are connected to their factor inputs via arrows. Such diagrams denote linear causality and illustrate simple chains of logic and associated arithmetic.

- Causal Loop Diagrams, such as seen in Figure 6, which depict circular feedback mechanisms. Causal loop diagrams can either counteract the direction of change (called a negative or balancing feedback loop) or accentuate it (known as a positive or reinforcing feedback loop).

- Stock-Flow Diagrams, such as seen in Figure 8, which depict some of the underlying system dynamics structures contained in the BSM (as rendered using the STELLA software package and its application-specific iconography).

Interested readers can find much more detailed information on causal loop diagrams in various online tutorials and in Chapter 5 of Business Dynamics: Systems thinking and modeling for a complex world (Sterman 2000).

More on stock-flow diagrams can be found on the website of the STELLA software developer (www.iseesystems.com) or in An Introduction to Systems Thinking with STELLA (Richmond 2004).

Diagrams in this report have been simplified/stylized to call out key aspects of the principle(s) depicted. For example, stock-flow diagrams do not show every element and/or variable contained in the actual STELLA model. Similarly, causal loop diagrams show key components and overall feedback loop polarity, but do not necessarily depict each portion of the underlying logic.

2

2 An Overview of the BSM Architecture

BSM has been designed in a top-down, modular fashion. In developing the model, we have taken care to create a structure that is transparent, modular, and extensible. This modular approach enables standalone analysis of individual model components as well as testing of different module combinations. As shown below in Figure 2, the model is framed as a set of interconnected sectors and modules.
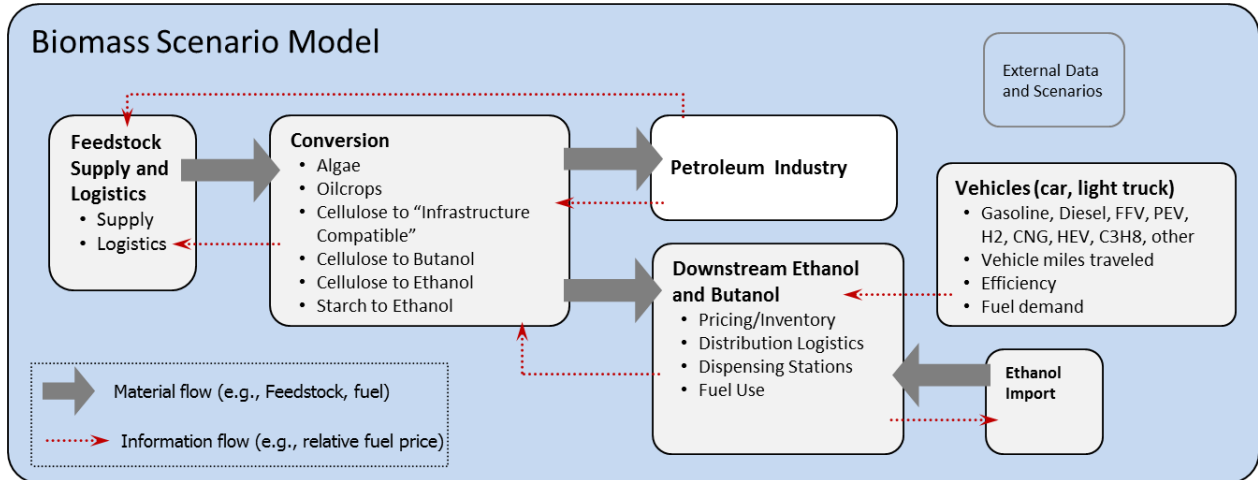
Figure 2. Overview of BSM structure

2.1 Feedstock Supply and Logistics

The feedstock supply and logistics sector captures the dynamics of cellulosic, oil crop, and starch feedstock supply from agricultural lands within the context of the operation of the United States agricultural system. It captures harvesting and transportation logistics associated with cellulosic feedstock. It also captures feedstock supply and logistics associated with both forest, urban, and agricultural residues.

Feedstock production from agricultural land occurs against the backdrop of other uses of the agricultural land base. These include commodity crop production (corn, wheat, soybean, small grains, cotton), hay, pasture, and Conservation Reserve Program (CRP) land. The agricultural production system disaggregated regionally into 10 production regions taken from the United States Department of Agriculture (USDA). These production regions are shown below, in Figure 3:

3

USDA Farm Production Regions
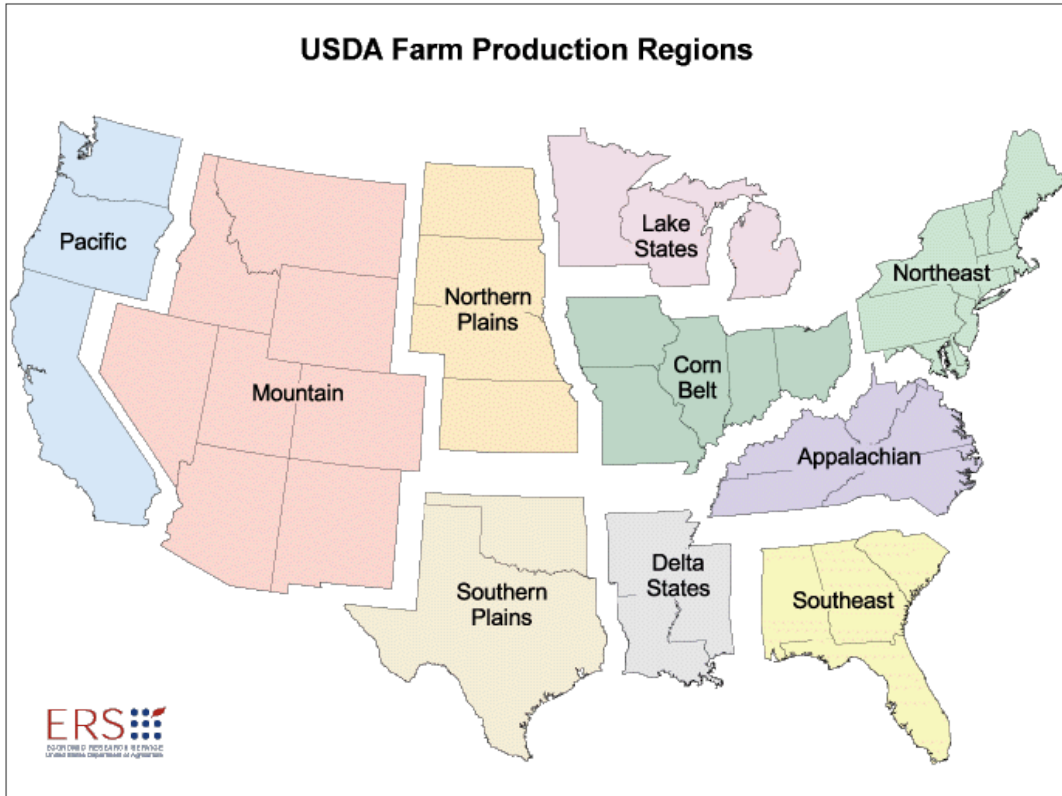
Figure 3. USDA farm production regions

2.2 Conversion

The conversion sector is composed of six different conversion modules, each corresponding to a different set of pathways for production of biofuels.

- Starch to Ethanol: This module represents the conversion capacity acquisition and utilization dynamics associated with the existing starch (corn) ethanol industry. This industry is considered to be mature; hence, the module provides a simple representation of the financial logic that controls acquisition and utilization of commercial scale corn ethanol facilities. This module is disaggregated by USDA production regions.

- Cellulose to Ethanol: This module captures the development of the cellulose-to-ethanol conversion industry. Biochemical and thermochemical conversion options are considered on a USDA-regionalized basis. The module represents pilot, demonstration, pioneer-commercial- and full-commercial-scale facilities. It includes learning curve dynamics, investment decision logic, and utilization logic for both pioneer and full commercial scale facilities.

- Cellulose to Butanol: This module captures the development of the cellulose-to-butanol conversion industry. In BSM, butanol serves as an industrial solvent and as a substitute for ethanol in gasoline blends. A single, regionally-disaggregated cellulose-to-butanol conversion option is captured in the model. The module represents pilot, demonstration, pioneer-commercial and full-commercial scale facilities. It includes learning curve

4

dynamics, investment decision logic, and utilization logic for both pioneer and full commercial scale facilities.

- Cellulose to "Refinery Ready/Infrastructure Compatable:" This module captures the industry development of cellulose-to-refinery-ready "infrastructure compatible" conversion processes. The model structure can accommodate these conversion options:

  o Fast Pyrolysis

  o Fischer-Tropsch

  o Methanol to Gasoline

  o Catalytic Pyrolysis

  o Fermentation

  o Aqueous Phase Reforming

  As with other cellulosic modules, this module is disaggregated by USDA regions. It provides a representation of pilot, demonstration, pioneer-commercial and full-commercial scale facilities. It includes learning curve dynamics, investment decision logic, and utilization logic for both pioneer and full commercial scale facilities. Multiple products or product substrates can be produced, including gasoline, diesel, and jet fuel. The "drop-in point" for various products is determined as a scenario variable.

- Oil crops: The Oil crop module captures development of conversion capacity for soy-to-refinery and "other" oilseed to refinery processes. Oil crop conversion facilities are represented as U.S. aggregates (rather than disaggregated by USDA production region). The module represents pilot, demonstration, pioneer-commercial and full-commercial scale facilities. It includes learning curve dynamics, investment decision logic, and utilization logic for both pioneer and full commercial scale facilities.

- Algae: The Algae model represents open pond, photobioreactor, and heterotrophic conversion options. It is not geographically disaggregated. Algae feedstock production is presumed to be vertically integrated in the algae to refinery-ready system. The module represents pilot, demonstration, pioneer-commercial and full-commercial scale facilities. It includes learning curve dynamics, investment decision logic, and utilization logic for both pioneer and full commercial scale facilities.

In addition to the six conversion modules, the conversion sector includes a simple module that knits together the "attractiveness" of the various investments in conversion options, allocating limited facility construction capacity among these options based on their perceived relative economic value.

2.3 Petroleum Industry

The petroleum industry sector comprises scenario inputs around crude oil prices, It provides logic that translates these prices into price inputs for the various refinery-ready conversion modules as well as the pricing/inventory module of the downstream ethanol/butanol sector. Additionally, the petroleum industry model provides accounting logic that captures displacement of crude by biofuel-derived infrastructure compatible fuels.

5

2.4 Downstream Ethanol and Butanol

The downstream ethanol and butanol sector is composed of a set of five modules. These modules capture ethanol and butanol activities "downstream" of conversion.

- Pricing/Inventory: This module captures pricing and inventory dynamics for both ethanol and bio-based butanol. Ethanol flows into two distinct but coupled markets: the "low-blend" gasoline market—10 to 15% by volume ethanol in petroleum blendstock—and the "high-blend" market associated with flexible fuel vehicles. Bio-butanol is assumed to serve as a substitute for ethanol in the low-blend market, and also can supplant butanol produced by other processes in the industrial market.

- Distribution logistics: This module provides a very simple representation of the regional build-out of the distribution network for fuel ethanol.

- Dispensing stations: The Dispensing Station Module addresses the regional acquisition of tankage and equipment capable of dispensing high ethanol blends into flex-fuel capable vehicles. Build-out of E85-capable stations is driven by economic considerations, and is constrained by regional availability of ethanol from the distribution network.

- Fuel use: The mix of low-ethanol-blend vs. high-ethanol-blend consumption is determined by the relative economics of the two products as constrained by the regional availability of ethanol for high-blend consumption through dispensing stations.

2.5 Vehicles

The vehicle scenario module functions primarily as an accounting structure, which is used in BSM to keep track of the cumulative effect of multiple scenarios around volume, vehicle mix, vehicle efficiency, and vehicle miles traveled for the car and light-duty truck sectors. Its structure captures acquisition, aging, and retirement of vehicles, as well as the translation of vehicles into potential demand for fuel.

2.6 Ethanol Import

The ethanol import module provides a simple representation of the evolution of non-domestic ethanol production capacity. It generates imports of ethanol into the United States based on a price differential as perceived from outside the U.S. This structure enables the model to capture historical patterns of growth and decline in imports of fuel ethanol. It is structured to facilitate exploration of multiple scenarios around production cost.

3 A More Detailed View of BSM

3.1 Feedstock Supply and Logistics

The feedstock supply and logistics sector is responsible for generating cellulosic, starch, and oil crop feedstocks for the conversion sector in BSM. The U.S. agricultural system forms the context for the production of a significant portion of these feedstocks. Accordingly, in developing the feedstock supply and logistics sector we have taken care to respect both the physical (land use) and economic aspects of U.S. agriculture. The sector is divided into two modules: Feedstock Supply and Feedstock Logistics.

Feedstock Supply

Feedstock supply refers to the production of different feedstocks required as substrate for conversion. Feedstocks generated by the Feedstock Supply Module are summarized in Table 1.

Table 1. Summary of Feedstocks Produced by Feedstock Supply Module

Feedstock	Source	Use	Notes
Corn	Crop Land	Ethanol	
Soy	Crop Land	"Refinery-Ready" fuels	
Other oil seed	Crop Land (small grains)	"Refinery-Ready" fuels	Model does not explicitly represent land allocation to other oil seed (e.g., rapeseed)
Crop Residue	Crop Land	Ethanol | Butanol | "Refinery-Ready" fuels	Model allows residue collection from corn, wheat, other grains, cotton
Herbaceous cellulosic energy crop	Crop land Pasture land	Ethanol | Butanol "Refinery-Ready" fuels	
Woody cellulosic energy crop	Crop land Pasture land	Ethanol | Butanol "Refinery-Ready" fuels	
Pasture	Pasture land	Ethanol | Butanol "Refinery-Ready" fuels	
Urban residue	Urban areas	Ethanol | Butanol "Refinery-Ready" fuels	Represented as simple price-response supply curve
Forest residue	Forest lands	Ethanol | Butanol "Refinery-Ready" fuels	Represented as simple price-response supply curve

As indicated in Table 1, urban and forest residue feedstocks are generated using simple price-supply relationships. All other feedstocks are produced by the agricultural land base. Figure 4 identifies the different land categories represented within the Feedstock Supply Module.

7
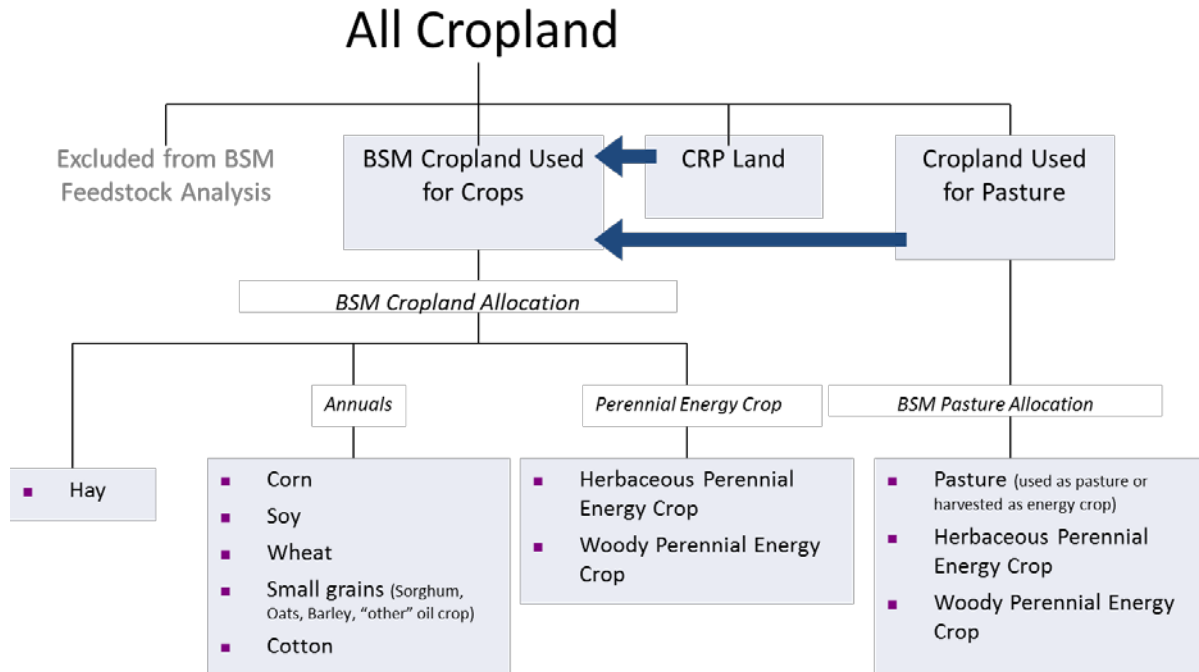
Figure 4. Land categories represented within Feedstock Supply Module

Within each of the ten USDA regions represented in the model, land is divided among three high-level categories: cropland used for crops, Conservation Reserve Program (CRP) land, and cropland used for pasture. These land bases are typically treated as static quantities over the course of a simulation run. However, as indicated in Figure 4, BSM structure supports scenarios that will cause land to move from CRP or pasture into cropland used for crops. Within each land base, land is allocated among different uses based on expected relative per-acre grower payment accruing to producers from the various products. Land allocation is region-specific, reflecting the production economics of different crops in different regions. Allocation of land to cellulosic crops is more restrictive: only those producers who have adopted the practice of producing cellulosic products (either residue or perennials) consider cellulosic grower payments in their decision making. "New practice" producers can grow over time based on the potential profitability of cellulosics, as constrained by the requirements of the existing and prospective conversion facilities.

Detail around the land allocation decision process is shown in Figure 5.

8
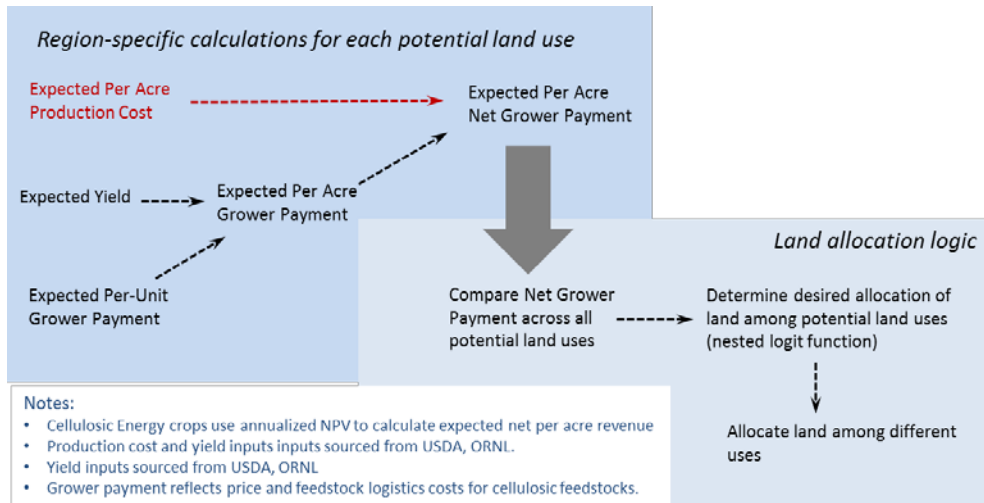
Figure 5. The agricultural land allocation algorithm

Not shown in Figure 5, but essential to the dynamics of BSM, is the logic surrounding pricing for the various commodity crops, cellulosic products, and hay. This logic is central to a feedback mechanism that uses land allocation to equilibrate production and consumption across all product categories in the model. A more detailed treatment of the pricing structure used in BSM is provided in Appendix A. Figure 6 shows in simple terms the feedbacks around price in the Feedstock Supply Module.
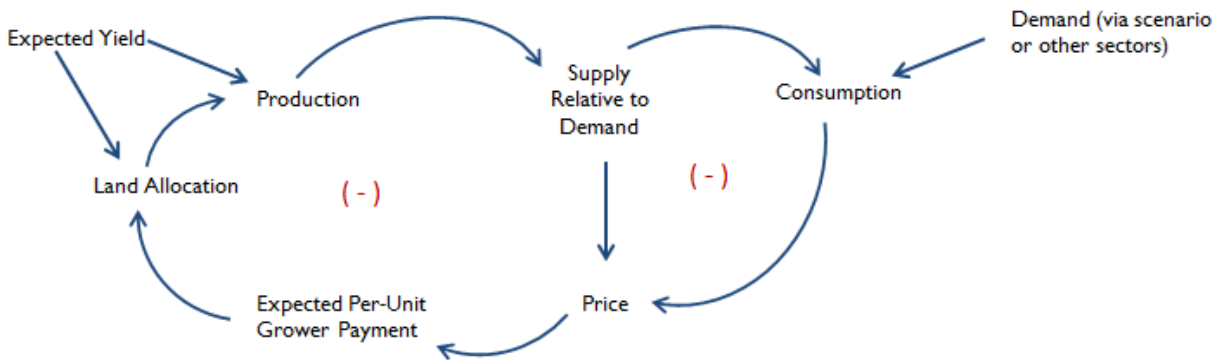
Figure 6. Price feedbacks in Feedstock Supply Module. (-) indicates negative feedback loop.

Feedstock Logistics

The Feedstock Logistics Module provides a simple accounting structure that captures the following costs:

- Harvesting and collection

- Transport from "farm gate" to "plant gate"

- Storage, queuing, handling, and pre-processing between farm gate and plant gate

9

These costs are used to translate the per-ton price of cellulosic feedstock at the plant gate into a per-ton grower payment at the farm gate. In developing the Feedstock Logistics Module, we have drawn from analyses of the Biomass Logistics Model (BLM) developed at the Idaho National Laboratory (INL). The Feedstock Logistics Module supports cost accounting for both pioneer and advanced storage, pre-processing, and queuing/handling processes.

The Feedstock Logistics Module underscores the high degree of interplay among different cost components. For example, truck transport is viewed as a primary mechanism for moving feedstock from farm gate to plant gate. Depending upon the feedstock involved, the mass transported on the truck varies, with residue resulting in significantly lighter loads than woody cellulosic crops. Other things equal, this implies a higher logistics cost per ton for residues than for woody cellulosic crops.

Additionally, the logistics module emphasizes the importance of the travel distance from farm to conversion facility. The model estimates farm-plant distances regionally by considering the following components:

- The total number of cellulosic plants requiring agriculturally-produced feedstock

- The total volume of agricultural land allocated to producing cellulosic feedstock

- The aggregate average yield of those producing acres

- An estimate of the fraction of land within the "plant-shed" that is available for cellulosic harvesting

- Geometric factors that relate the resultant plant-shed area to average travel distance from farm to plant.

3.2 Conversion Sector

The conversion sector is responsible for transforming feedstock into liquid fuels, including ethanol, butanol, and refinery-ready fuels (gasoline, diesel, jet fuel) suitable for insertion into the existing fuel infrastructure as refinery feedstocks, blendstocks, or finished products. In BSM, the conversion module comprises a significant fraction of the overall model structure. It consists of seven modules. Six of these modules look at the dynamics of industry development for sets of conversion pathways. These dynamics include operations at different scale factors, learning along multiple dimensions, logic surrounding the attractiveness of investment in new facilities, and utilization of existing facilities. A seventh module compares investment attractiveness across all conversion options, allocating scarce investment capacity among these options based on their relative net present value.

Table 2 provides an overview of the scope of the conversion sector. As indicated in Table 2, there is significant overlap among the different industry development modules. In particular, most modules share the following characteristics:

- Multiple conversion options, represented using an arrayed variable structure

- Regional disaggregation, following the Feedstock Supply Module's use of ten USDA production regions

10

- Incorporation of pre-commercial pilot and demo-scale operations
- Representation of pioneer-commercial-scale operations
- Representation of full-scale operations
- Learning curve dynamics.

Table 2. "Dimensionality" of Conversion Sector

	Conversion Options	Regional?	Feedstock	Products	Pilot & Demo Ops?	Pioneer-scale Ops?	Full-scale Ops	Learning Curve Dynamics?
Starch to Ethanol	Single pathway	Yes	Corn	Ethanol	No	No	Yes	No (assume mature industry)
Cellulose to Ethanol	Biochemical Thermochemical	Yes	Cellulosic Feedstock	Ethanol	Yes	Yes	Yes	Yes
Cellulose to Butanol	Single Pathway	Yes	Cellulosic Feedstock	Butanol	Yes	Yes	Yes	Yes
Cellulose to Refinery	Fast Pyrolysis Fischer-Tropsch Methanol to Gasoline Catalytic Pyrolysis Fermentation APR	Yes	Cellulosic Feedstock	Gasoline Diesel Jet fuel (3 drop-in points)	Yes	Yes	Yes	Yes
Oil Crop to Refinery	Soy Other	No	Oil crop	Diesel Jet fuel (3 drop-in points)	Yes	Yes	Yes	Yes
Algae to Refinery	Pond Photobioreactor Heterotrophic	No	Algae— treated as part of conversion process	Diesel Jet fuel (3 drop-in points)	Yes	Yes	Yes	Yes (feedstock supply considered endogenous to module and subject to learning curve)

Within the sector, there are a few departures from the generic structure. For example, in the oil crop and algae modules, regional production of feedstock is of secondary importance. For this reason the BSM does not disaggregate these modules by region. The starch to ethanol industry, to take another example, is assumed to have reached maturity. Because the industry is mature in

the model, there is no need to represent the dynamics of pilot, demo, or pioneer scale operations. Nor is there a requirement to represent learning curve dynamics. In the algal module, feedstock production is considered as endogenous to the algae system rather than produced by the feedstock supply sector. Algal feedstock production costs are subject to learning curves in the algae module.

Within the typical conversion module, there are multiple processes that govern the development of the conversion options under consideration and their production of fuel. These processes are centered on:

- Pilot and demonstration scale operations
- Pioneer-commercial-scale operations
- Full-commercial-scale operations
- Expected economic value of the "next" investment
- Allocation of scarce capital in the investment decision
- Learning along multiple dimensions
- Industry aggregate average utilization rates for existing facilities.

Figure 7 provides a schematic representation of key interconnections among these processes.
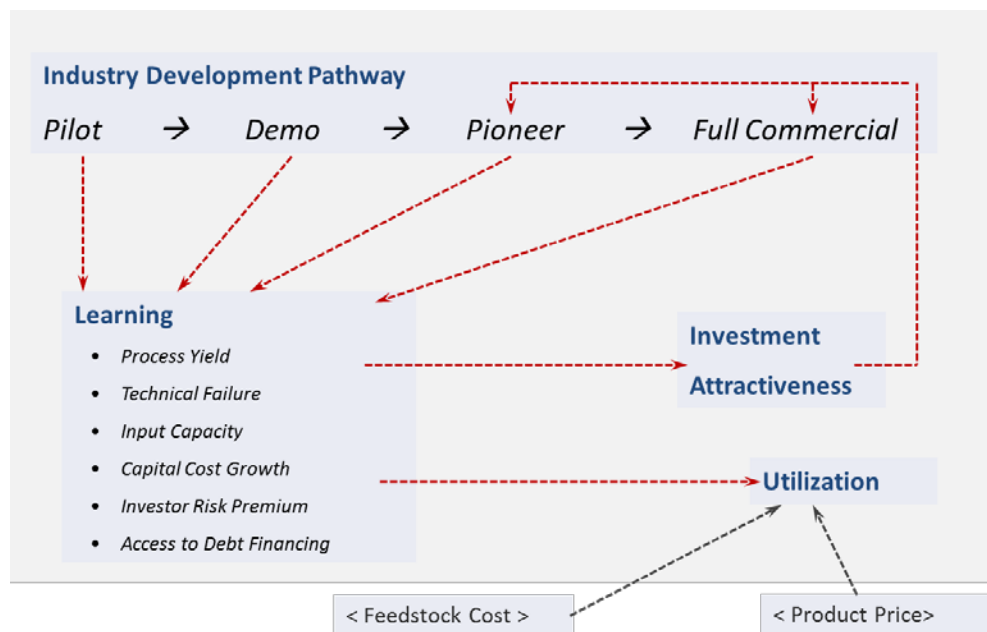
Figure 7. Key interactions within the typical conversion module

Pilot- and Demonstration-Scale Operations

Pilot and demonstration-scale operations are represented simply in the model. Figure 8 shows the stock/flow structure of these operations.

12
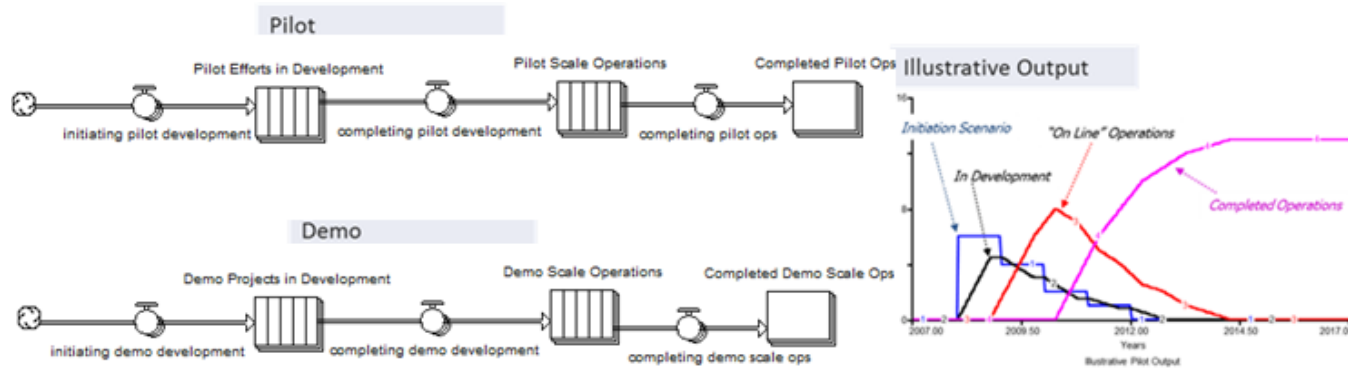
Figure 8. Pre-commercial structure and illustrative output

There are several important features of the pilot and demonstration scale structures. First, the structure for each pre-commercial-scale operation is arrayed, based on the number of conversion options at play within the module. Technology switches, set by the user, enable the activation or deactivation of each conversion option. Second, both pilot and demo operations are specified as exogenous scenario inputs. These scenario inputs enable an arbitrary pattern of initiation to be specified by the end user as a scenario. Third, the model explicitly represents the dwell time between initiation and completion of development using conveyors. Fourth, the time for which operations are active—used in the model to generate learning—is limited in duration. Finally, note that cumulative completed operations are tracked by the structure. Figure 8 provides illustrative output that translates an arbitrary initiation scenario into development, "on line," and completed operations.

Pioneer-Commercial-Scale Operations

The model accommodates two commercial scale operations: Pioneer-and full-commercial scale operations. In the model, pioneer-commercial-scale facilities are often the first commercial scale plants to come on line. These facilities have a smaller (about one-third) capacity than full-commercial-scale facilities. They do not take full advantage of economies of scale. Hence, in typical simulations of BSM, subsidies are required to stimulate investment in pioneer plants. Note that the starch to ethanol module excludes pioneer facilities from analysis.

As shown in Figure 9, two stock/flow chains are used to account for pioneer scale plants. Depending on the module in question, these chains are arrayed by conversion option and/or by region (Refer to Table 2, above, for details on the "dimensionality" of each conversion module.) The top chain represents the *number* of plants in design and construction, in startup, and in use. The bottom chain is a co-flow structure that is used to account for the *process yield* (gallons of output per ton of feedstock input.) These two concepts—facilities and process yield—jointly determine the output capacity for pioneer facilities in the aggregate. Output capacity is a reflection of the total ability of pioneer-scale facilities to produce fuel via a particular conversion pathway.

The co-flow structure is essential for the accurate accounting of facilities and their associated process yields. Whenever a new facility enters the system through the initiating flow, the model samples the current state of the industry process yield for the associated conversion option. This process yield then moves along with the facility through the development process, eventually

13

being used as an input for the average process yield of on-line facilities. This structure enables the model to dynamically track the cumulative impact of growth in process yields for "new" plants, as the industry moves from "blue sky" to "nth plant maturity."
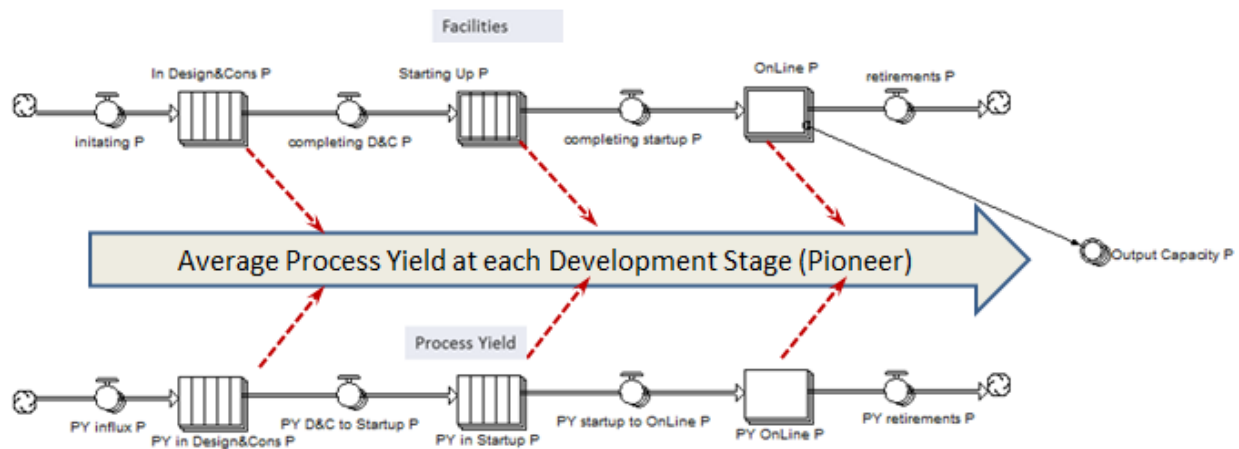
Figure 9. Accounting for pioneer facilities, process yield, and output capacity

Note that structure has been provided to account for retirement of plants. In the October 2012 version of the model, the assumed retirement fraction is zero, which implies that no plants are taken permanently off line over the course of a simulation. The logic that controls utilization factors, discussed below, accounts for the dynamics of short-term plant idling in response to market forces.

Full-Commercial-Scale Operations

Figure 10 shows the structure that accounts for full-commercial-scale operations in the model. As with pioneer plants, commercial plants use two stock/flow chains to represent the design and construction, the start-up, and the on-line phases of the facility life cycle. As with pioneer plants, the structure for commercial operations is arrayed by conversion option and/or by region within each module.

A comparison of Figure 9 and Figure 10 will reveal two notable differences between the pioneer and commercial accounting structures. First, note that in contrast to pioneer facilities, commercial facilities in the start-up phase are assumed to contribute to the overall output capacity. A utilization rate (less than 1) is assumed for plants during the period that they are in startup.

Second, note that a new flow has made its way into the process yield chain for commercial scale facilities. This flow enables the model to capture the effect of process yield improvements to be incorporated into the existing capital stock. The user of the model can specify the specific rate at which a yield gap—measured as the discrepancy between the state of the industry process yield for a particular conversion option and the existing industry average process yield for that conversion option—is eliminated.

14
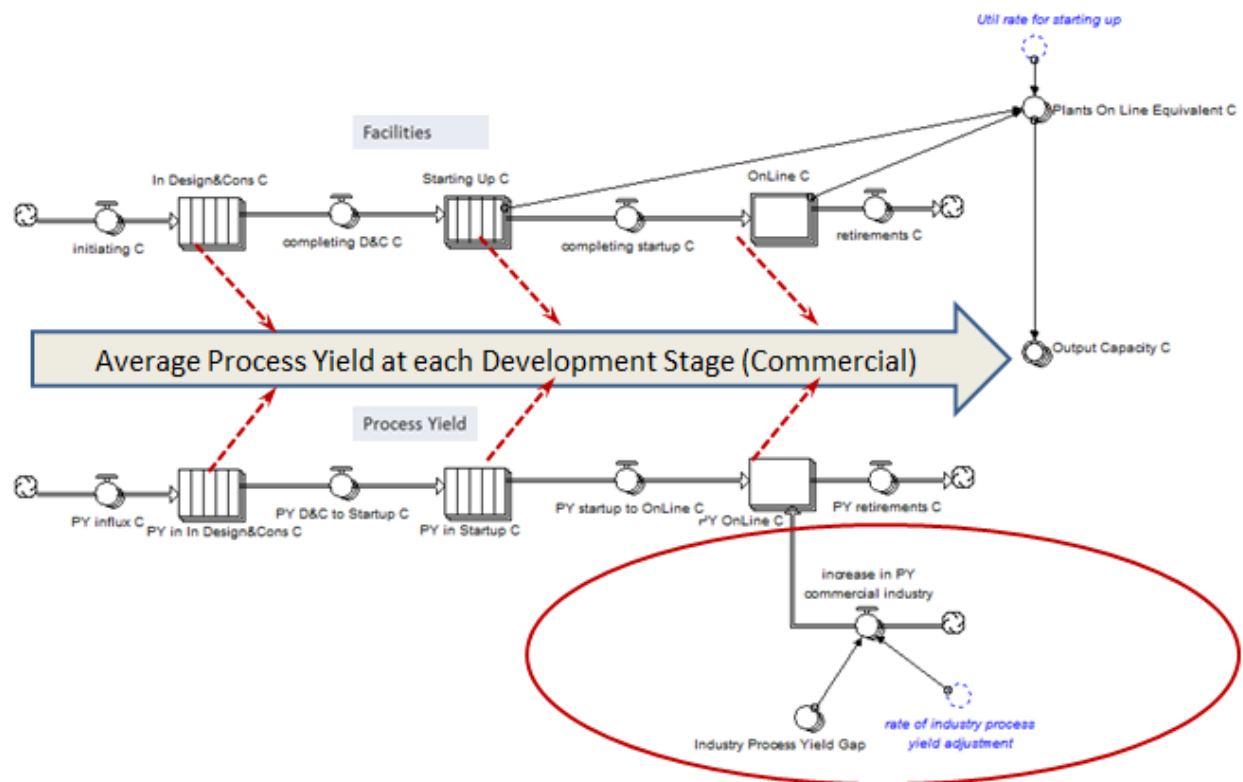
Figure 10. Accounting for commercial operations, process yield, and output capacity

Expected Economic Value of the "Next" Investment

In developing BSM, it was essential to develop a simple, defensible mechanism for determining the viability of investment in the "next" plant, at either pioneer or commercial scale, for the various conversion options within the regions under consideration. In contrast to a static calculation, we needed a dynamic economic mechanism to facilitate industry growth. In response to this challenge, we developed a structure culminating in a Net Present Value (NPV) calculation. At any point in simulated time, this structure captures important streams of costs and revenues associated with a prospective project investment. By discounting these streams to the present, it captures the dynamics of an evolving industry using a simple metric that enables comparison of prospective investments across multiple conversion options, regions, and scales. In turn, this metric enables the model to allocate scarce capital toward its highest valued uses.

A parallel algorithm is used for NPV calculations within each conversion module. As appropriate to each module, the algorithm reflects conversion options, regional considerations, and scale. Wherever possible, the algorithm operates at the highest possible degree of aggregation by rolling up sub-categories into high-level summaries. For example, for purposes of the NPV calculation, factor inputs and expected per-gallon revenues are held constant over the plant lifetime. Figure 11 provides a simple influence diagram showing the logic flow leading to the NPV calculation.

Figure 11. Logic of NPV calculation

In developing the NPV logic, we have adopted some important simplifications. In addition to simplifications around revenue streams, we assume straight-line depreciation of the plant in question. This significantly reduces detail complexity in the model. Additionally, we divide the overall project life cycle into distinct phases, as shown in Figure 12. In the model, NPV calculations are made for each phase of the project life cycle, and then rolled up to create an overall NPV for the plant.

Figure 12. Phases of project life cycle

Allocation of Scarce Capital

Within the BSM conversion sector, then, multiple opportunities present themselves to potential investors at any point in time. At the extreme, thirteen conversion options can be active. Nine options compete across ten regions. Twelve conversion options exist at both pioneer and commercial scale. That's a lot of options for investors to consider! Within the model, each conversion module uses NPV as a basis for determining the attractiveness of the various investment options under consideration. This is done using a logit function. (More on the use of logit functions in BSM can be found in Appendix B.) The resultant attractiveness metrics are then compared within the Relative Attractiveness (RAT) module, which also includes a default "other" investment category. The relative attractiveness for each alternative is then applied to a scenario-driven maximum construction capacity, which generates a platform and scale-specific

16

yearly start rate. This "desired" start flux is communicated back into the conversion modules, where it is allocated regionally if required and "batchified" so as to send a discrete signal to begin plant development. Figure 13 provides a view into the logic surrounding allocation of capital.
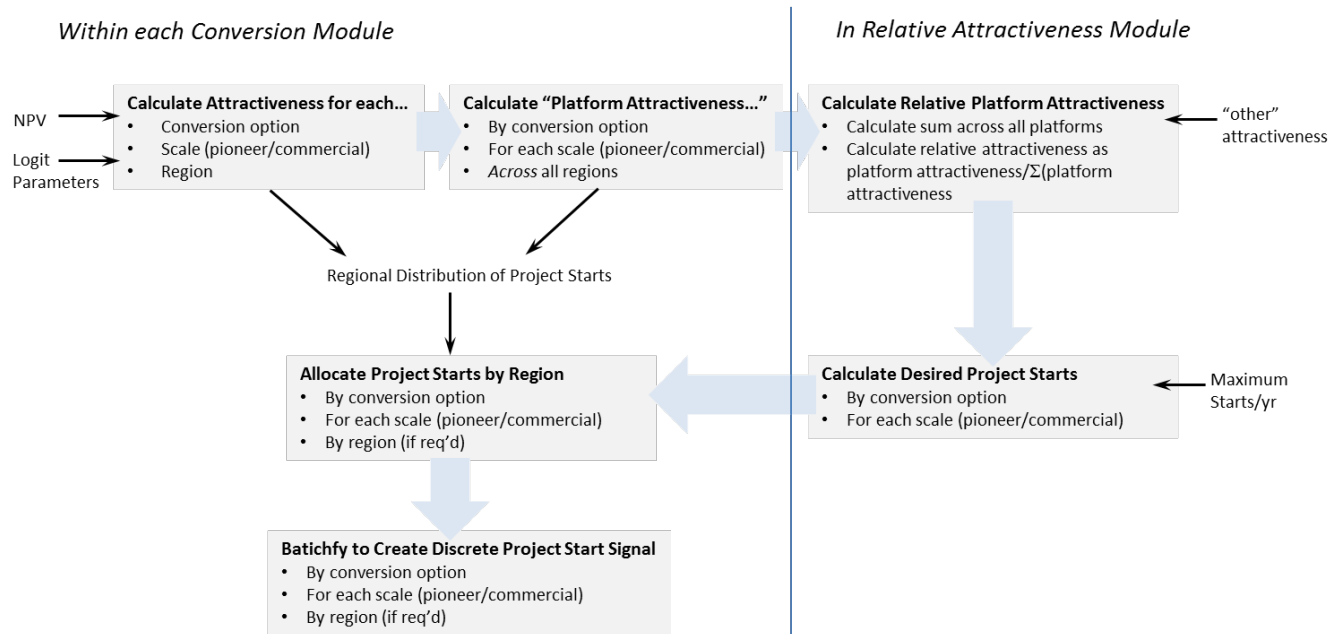
Figure 13. Translating NPV into project start signal

Learning along Multiple Dimensions

For most conversion options under consideration in BSM, the initial performance along multiple dimensions would fall far short of expected mature industry (or "Nth plant") performance, as reflected NREL and other design studies. Industry evolution is in no small measure the story of performance improvement that results from learning by doing. Merrow's research on cost growth in capital-intensive industries (Merrow 1983), for example, underscores the important role of experience at prior scale in reducing the risk of capital cost growth at commercial scale. Henderson's work with the Boston Consulting Group in the 1970s (Hax and Majluf 1982) demonstrates the role of learning as a cost-reduction strategy at commercial scales.

Given the important connections between learning and industry evolution, we needed to develop a simple, consistent, and defensible mechanism to translate the accumulation of experience into a set of performance parameters to represent the current "state of the industry" for each conversion option. Our approach, which we call "cascading learning curves," draws upon simple learning curve principles in order to address learning for multiple conversion options at multiple development stages, addressing multiple performance attributes. There are three fundamental tasks involved in the cascading learning curve approach:

- Develop separate cascading curves for each conversion option. By providing separate structure for each conversion option, we have created the possibility to separately characterize different initial conditions, mature industry conditions, and learning rates on a conversion option-specific basis.

17

- Capture learning for each conversion option at three distinct development stages. In BSM, we look at learning for pilot scale operations, for demonstration scale operations, and for commercial scale operations (including pilot and full commercial scale). A staged approach to learning enables us to capture prior scale effects (important for capital cost growth). It also enables us to explore the implications of stage-specific progress rates as well as the analysis of timing and placement of policy initiatives.

- Use learning to create indices of maturity. These indices of maturity, in turn, drive essential technology attributes that are used within BSM. Key attributes of performance for each conversion option are:

    o Process yield

    o Likelihood of "technical failure"

    o Feedstock throughput capacity—the degree to which facilities are able to perform at nameplate capacity

    o Capital cost growth—the premium in capital cost, beyond Nth plant estimate, which would be observed if development of a facility was begun today

    o Investor risk premium—the additional premium, beyond normal hurdle rate, that investors would require for investment in the facility

    o Access to debt financing—the portion of the expected facility capital cost that would be financed via borrowing (vs. equity investment).

The generic stock/flow learning curve structure is shown below in Figure 14. Annotations describe key aspects of the structure.
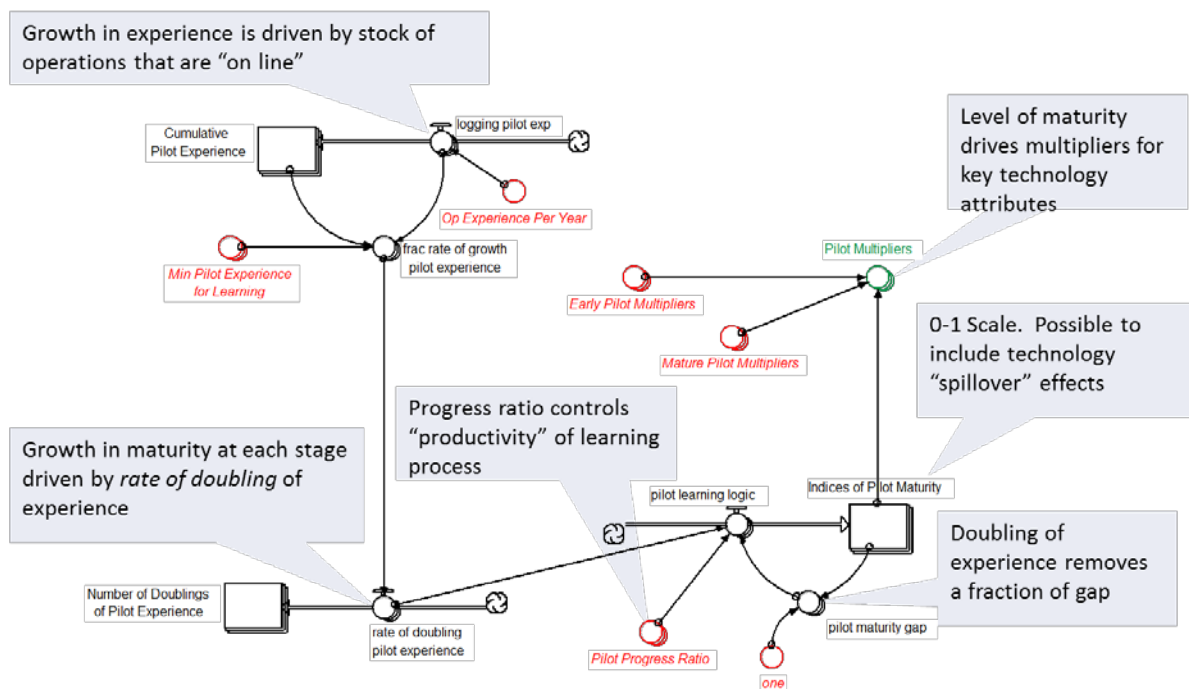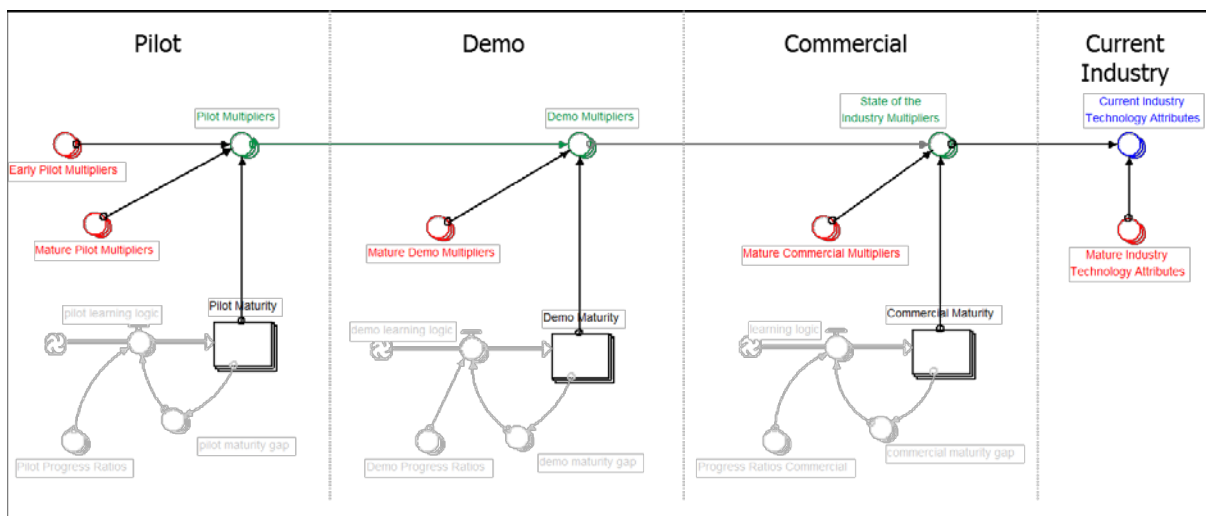

Figure 14. Industry learning curve structure

18

There are some differences between the representation of the learning curve structure in BSM and other, perhaps more common, formulations of learning curves. In the classic formulation of a learning curve, for example, a power law is used to relate cumulative experience to a single attribute such as cost. The asymptote of cost is often implicitly set to zero. By contrast, in BSM, cumulative learning at each stage of development is reflected along a 0-1 scale in pilot, demo, or commercial maturity. As experience accrues, the model calculates explicitly the rate at which experience is doubling. This rate of doubling is applied to a maturity gap (simply the difference between current maturity and full maturity) to generate learning. Maturity, in turn, drives movement along a vector of attributes.

A second set of differences involves the development stages over which learning is applied. While a typical learning curve analysis might consider cost reductions for relatively stable developed industries, in BSM we consider multiple attributes over multiple development stages. Figure 15 shows how the learning curves cascade over these development stages.

Figure 15. Cascading learning curves

At any point in simulated time, the current industry technology attributes reflect the performance and cost characteristics associated with an investment in a pioneer or full-commercial scale facility for a given conversion option. At each stage, multipliers that are passed on to the next stage are calculated as a weighted average, with the maturity level used as the weighting factor.

Dynamically, this structure enables BSM to jump from one performance trajectory to another based on the behavior of pilot, demo and commercial operations, as Figure 16 illustrates. In the figure, simple exogenously-defined scenarios for pilot, demo, and commercial scale operations drive learning at each stage. Cost and yield parameters follow three distinct pathways as the industry evolves.
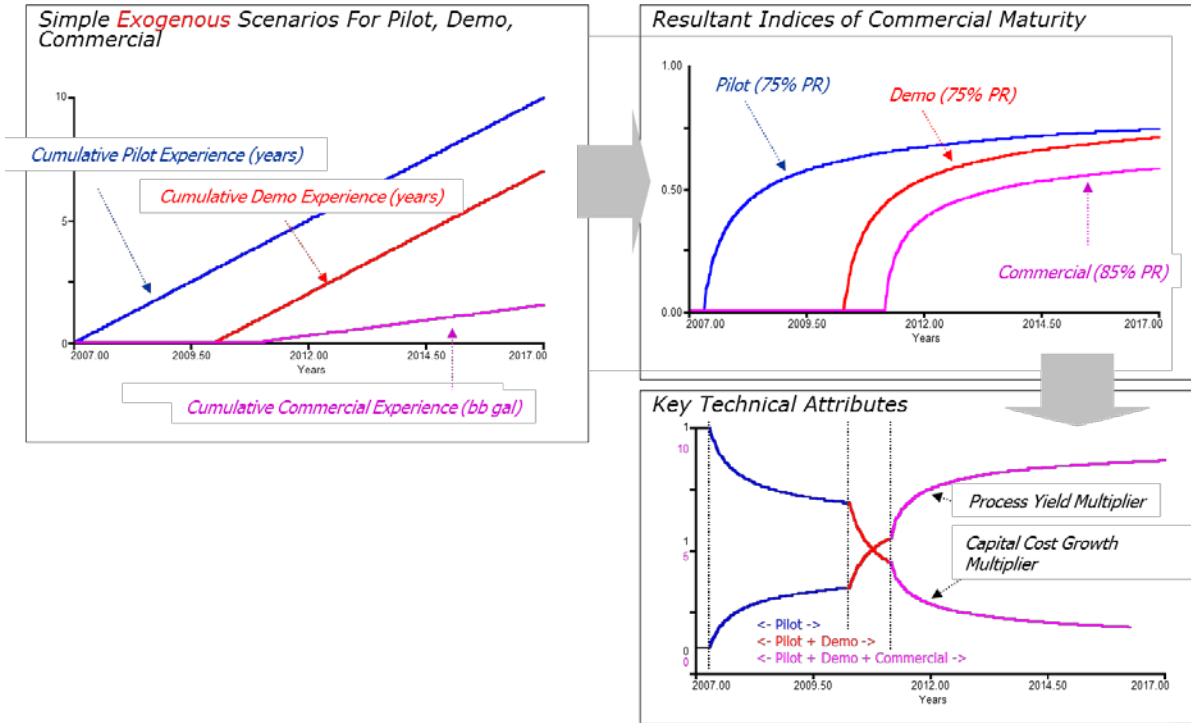
Figure 16. Illustrative learning curve dynamics

Learning curve dynamics, of course, do not occur in isolation from the overall dynamics of the industry. For a given conversion option, learning curves are at the heart of feedbacks that surround the investment process, and which can underwrite industry "take-off." These feedbacks are shown in a simple loop diagram in Figure 17. Each feedback is a positive feedback, in the sense that it tends to reinforce development of the conversion option in question.
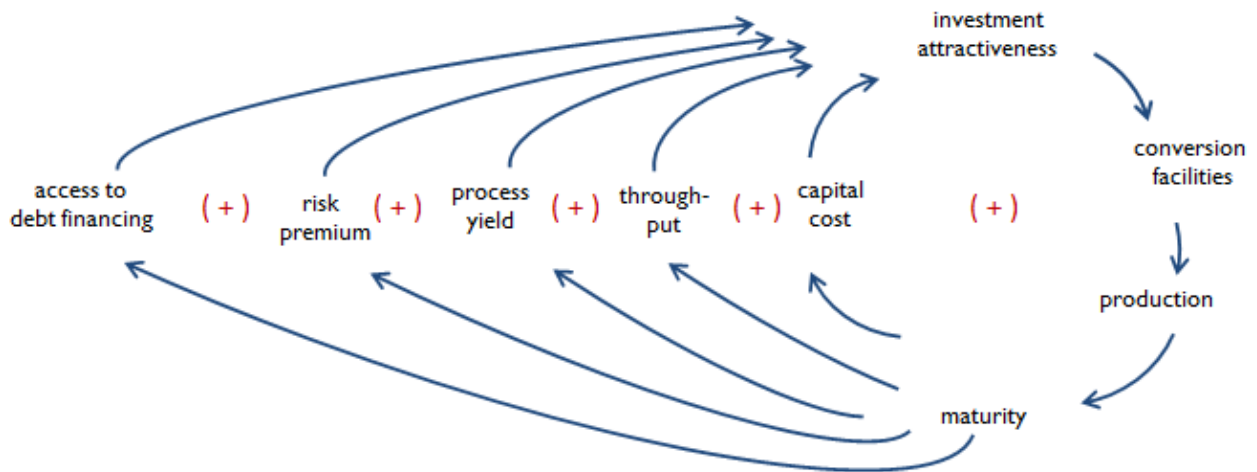
Figure 17. Key feedbacks in learning curve structure. (+) indicates positive feedback loop.

20

Utilization of Existing Facilities

Multiple processes are at work in the conversion sector to generate the production of biofuels. A final set of processes concerns the utilization of existing facilities. A fundamental premise of basic economics is that "sunk costs" don't matter. In BSM, conversion facilities are assumed to follow this premise; the capacity utilization rate for each conversion option (within each region, as appropriate) at either pioneer or full commercial scale is developed as a response to the "cost-price ratio" for its products. As the price received for its product (including any subsidies) grows relative to the per-gallon cost of producing that product (after factoring net per-gallon co-product revenues into the mix), utilization increases to its maximum. On the other hand, as the price-cost ratio declines below unity, utilization rates decline. The logic determining cost-price ratios and utilization is shown below in Figure 18.
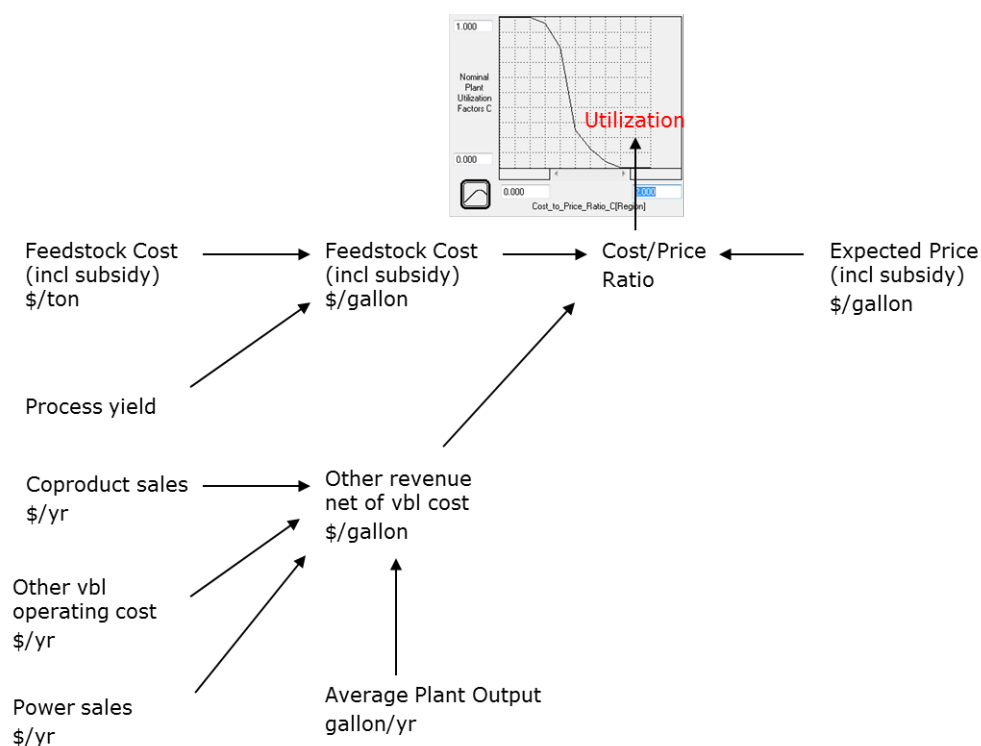
Figure 18. Determining utilization from cost-price ratios

Utilization, by controlling both production of products and consumption of feedstock, is central to the feedback structure within BSM.

Conversion Sector in Summary

Each conversion module within the conversion sector is built up from multiple simpler structures that represent pre-commercial demonstration and pilot scale operations, pioneer and full commercial scale operations, the expected economic value of investment, learning and utilization. These structures are connected within each module in order to generate products (diesel, jet fuel, gasoline, butanol, ethanol). They are connected across modules via the logic within the relative attractiveness module that allocates scarce investment capital. The conversion sector is connected upstream in the supply chain to the agricultural system through feedstock

21

supply dynamics, and to both the oil industry (algae, oil crop and cellulose to refinery-ready modules) and the downstream ethanol sector. These downstream sectors determine price (and in the case of ethanol, demand) signals which are sent to the conversion sector modules.

3.3 Downstream Ethanol Sector

The downstream ethanol (EtOH) sector comprises a set of interconnected modules that take fuel ethanol from conversion facilities to end users, both in low-blend (E10 or E15) and high-blend (nominally, E85) form. Additionally, the downstream sector contains logic that controls the use of bio-butanol as a substitute for ethanol in the low-blend market. The model assumes that physical characteristics of ethanol require separate infrastructure for distribution and dispensing than for petroleum-based fuels. A significant portion of the downstream sector, therefore, is focused on distribution and dispensing station dynamics.

Figure 19 provides a picture of the content of the downstream sector. As suggested by the diagram, downstream dynamics focus on the build-out of distribution infrastructure, the development of dispensing infrastructure, and decision making around fuel usage.
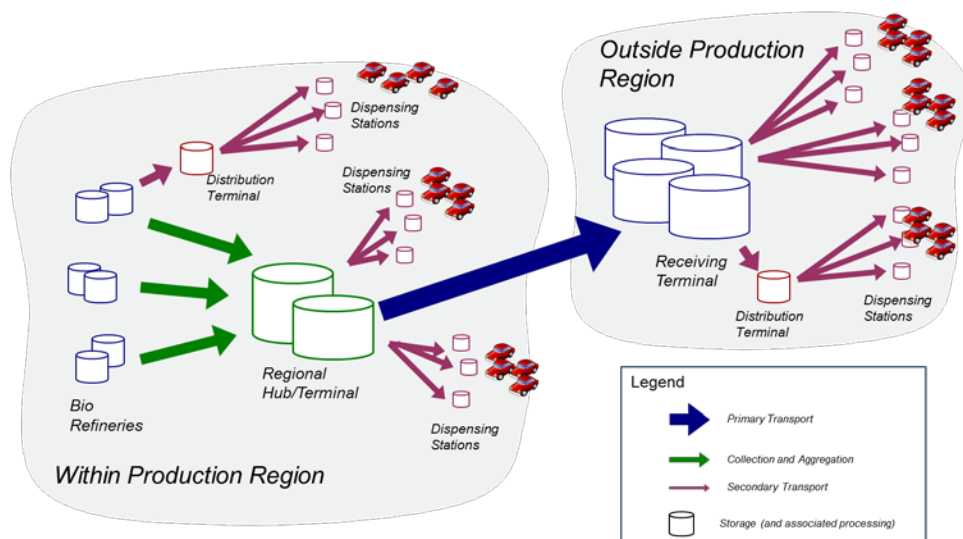
Figure 19. An overview of downstream dynamics

To support these dynamics, multiple modules comprise the downstream sector of BSM. These include:
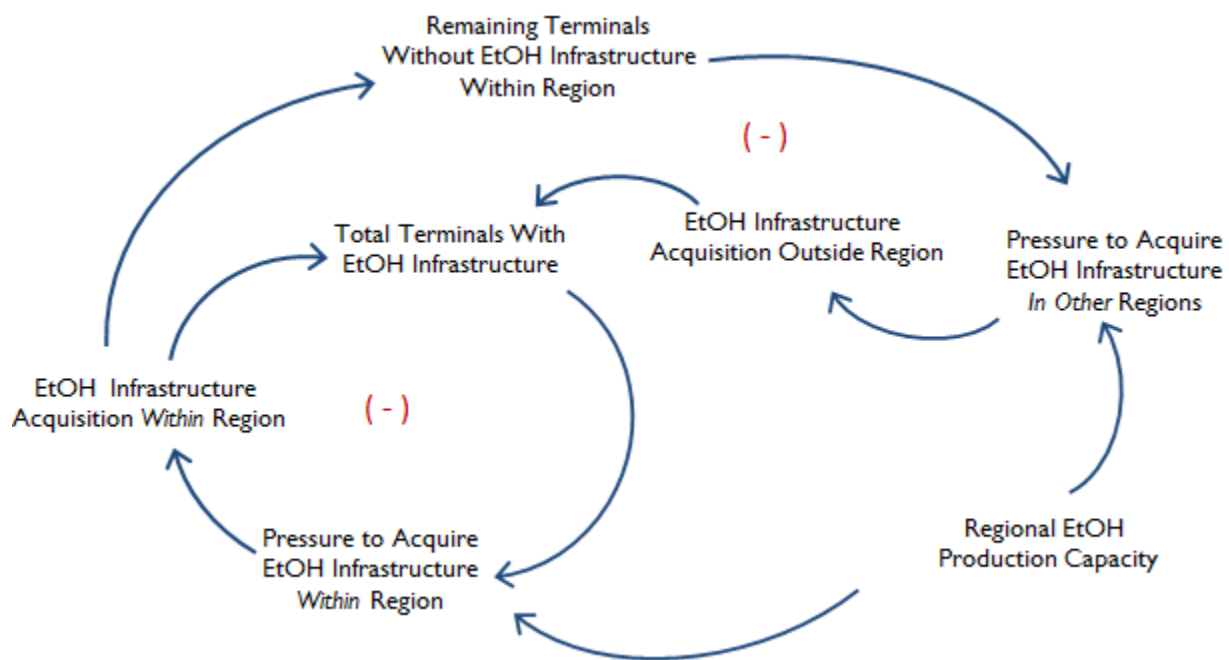
- Distribution logistics
- Dispensing station
- Fuel use
- Pricing and inventory

This report will provide a brief view into each module. Detailed analysis of downstream ethanol dynamics can be found in (Vimmerstedt, Bush, and Peterson 2012).

Distribution Logistics Module

One of the fundamental challenges associated with ethanol as a transportation fuel is its apparent incompatibility with existing infrastructure. The Distribution Logistics Module provides a very simple representation of the build-out of ethanol-friendly distribution infrastructure. Rather than speculating on the build-out of specific distribution modalities for ethanol (such as rail, barge, or dedicated pipeline), the Logistics Module focuses on capturing the implications of build-out on the rest of the downstream system. The structure focuses on the acquisition of ethanol infrastructure for terminals within each region. The module is silent on the specific details of infrastructure, instead focusing on the drivers and time delays associated regional build-out.

Figure 20 provides a simplified view of the logic that drives build-out within the Distribution Logistics Module.

Figure 20. Key feedbacks in distribution logistics structure. (-) indicates negative feedback loop.

A two-stage supply-push approach is embedded within the module. This supply push works first within a region, and then across regions. Within a region, the model seeks to balance ethanol production capacity against terminal capacity to distribute that ethanol. As production capacity within a region grows, there is pressure within the region for terminals to acquire ethanol-compatible distribution infrastructure.

Second, as build-out occurs within each region, any excess regional production capacity creates pressure for acquisition of infrastructure in other regions, in proportion to the terminal density within each region.

The result of this two-stage supply-push algorithm is an initial build-out of distribution infrastructure in ethanol producing regions, followed by a slower build-out in non-producing regions. Infrastructure coverage within any region constrains regional investment in ethanol

This report is available at no cost from the National Renewable Energy Laboratory (NREL) at www.nrel.gov/publications.

dispensing tankage and equipment, thus setting a limit on the uptake of ethanol in high-blend form.

Dispensing Station Module

The Dispensing Station Module focuses on the decision making associated with the acquisition and use of high-blend tankage and equipment by retail dispensing stations. The module considers roughly 120,000 stations, distributed both regionally and by ownership among oil-owned, branded independents, unbranded independents, and hypermart. The fundamental decision for each station is the acquisition of tankage and dispensing equipment required to dispense high-ethanol blends into flex-fuel vehicles (FFV). The module assumes that a small fraction (less than ten percent) of stations have repurposable mid-grade tanks. The capital cost of repurposing is assumed to be significantly lower than investment in new tankage and equipment for hi-blends ($20,000 vs. $60,000).

The basic logic within the Dispensing Station Module combines the physics of high-blend availability with the economics of the investment decision. Stations will not consider investment unless distribution infrastructure is sufficient within the region. They will not invest unless the investment makes economic sense, as reflected in a net present value calculation that captures the discounted stream of expected costs and benefits from the investment.

Thus, two fundamental structures are at play within the Dispensing Station Module. The first is an accounting structure that considers the movement of stations as they adopt high-blend tankage and equipment. The second provides a detailed view into the net present value (NPV) calculation that undergirds the decision to invest in high-blend tankage and equipment. Station movement structure is shown in Figure 21; NPV logic is shown in Figure 22.
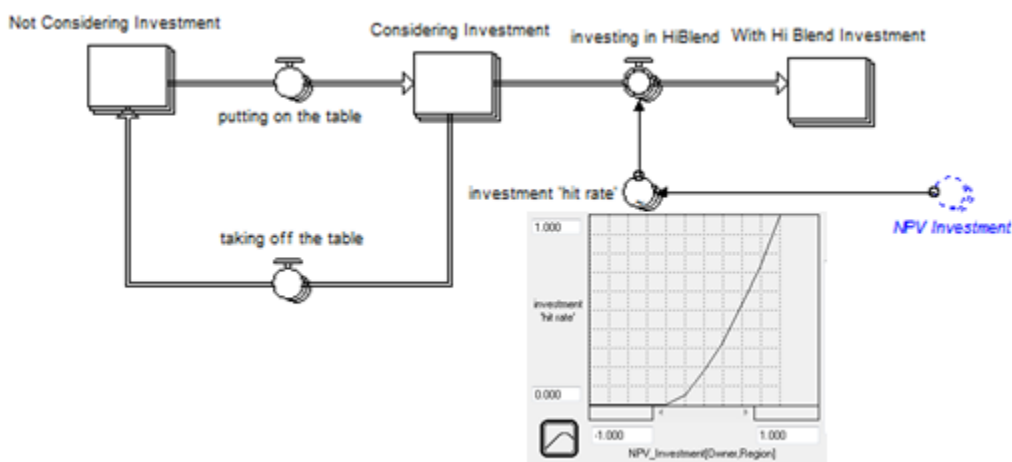
Figure 21. Dispensing station accounting structure. NPV calculation captures estimated costs and revenues of prospective investment

As shown in Figure 21, stations exist in one of three states with respect to investment in high-blend tankage and equipment. Depending on the dynamics of regional distribution infrastructure availability, each year a portion of those stations not considering investment transition to considering. Based on the economic viability of the investment, as reflected in the NPV of the

24

decision, the consideration of investment culminates in a decision to invest or to stop considering the decision. This investment process is disaggregated by region (so as to account for differential degrees of distribution infrastructure within each region), by ownership (to enable different potential affinities for high-blend ethanol sales among different ownership types, and to account for different business details for different ownership types), and by repurposed versus new investment (to account for different capital costs associated with repurposing vs. new investment in tankage and equipment).
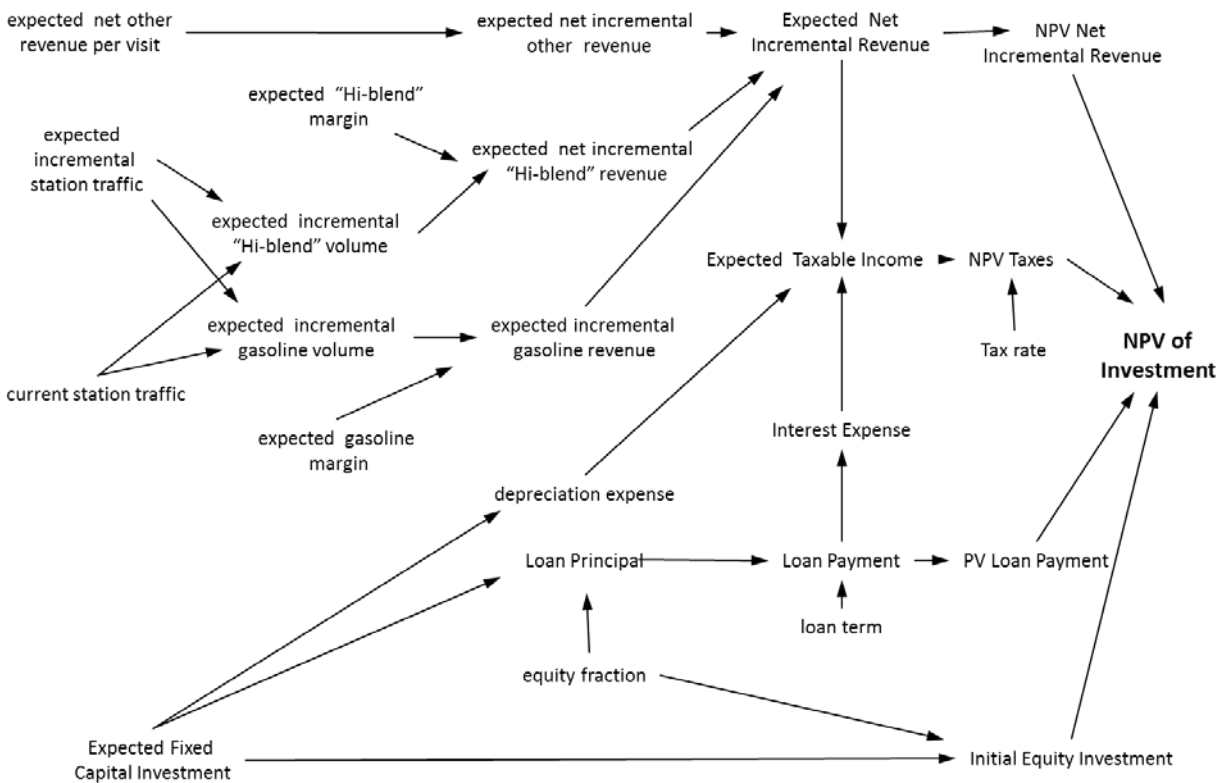
Figure 22. Logic behind NPV calculation for stations

As shown in Figure 22, the NPV calculation considers major categories of revenue and expense associated with station investment. In addition to the capital cost of the investment, the NPV calculation considers marginal cost and revenue streams associated with changes in the mix of high-blend versus "straight" gasoline sales, changes to station traffic (to account for first-mover advantage) and other revenues from c-store operations.

Just as the Distribution Logistics Module provides a context that constrains the acquisition of tankage and equipment for stations, the Dispensing Station Module provides a context for fuel use. Accessibility of high-blend stations within a region will constrain the potential for FFV to access high-blend fuels. Regional dispensing station coverage thus sets a physical limit on ethanol uptake in the system.

Fuel Use Module

The Fuel Use Module captures the both the effects of regional high-blend fuel availability and the effects of relative gasoline/high-blend pricing on the decision making for FFV owners, with respect to the use of high-ethanol fuel blends. The module contains two major interconnected components. The first component accounts for the affinity of FFV owners toward high-blend fuels. The second uses a logit function to allocate fuel use between for FFV owners who are "occasional" and "regular" users of high-blend fuels.

Figure 23 shows the accounting structure for FFV owners within a region.
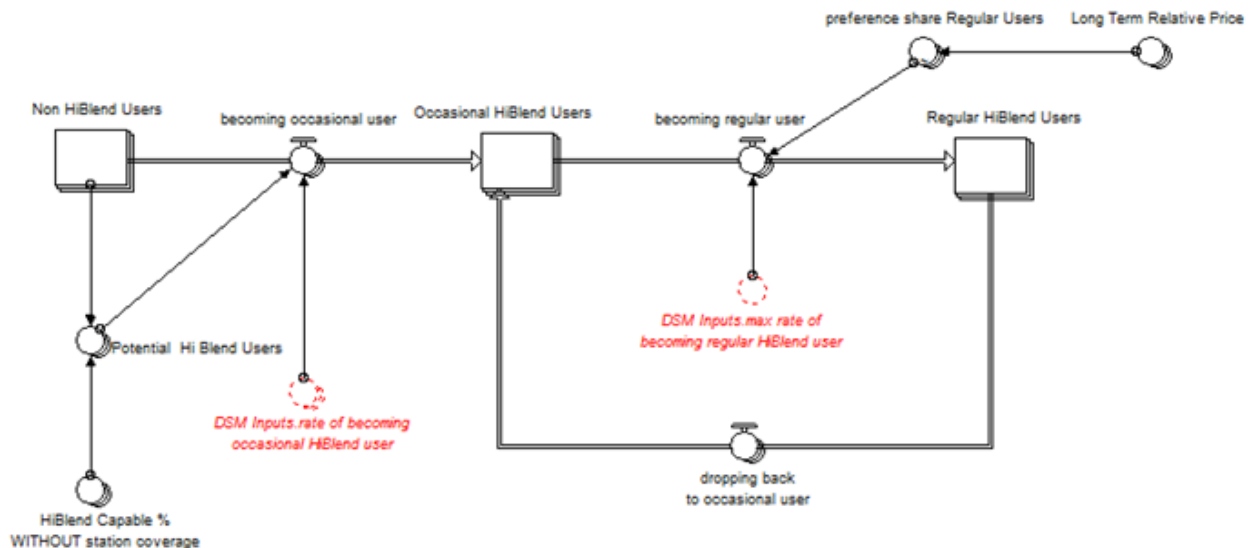
Figure 23. Flex-fuel vehicle accounting structure

As shown in Figure 23, FFV users (expressed as a % of regional FFV vehicles) are divided into three distinct categories: *non*-high-blend users, *occasional* high-blend users, and *regular* high-blend users. Non high-blend users do not use high-blend because a) they do not have access to stations that dispense high-blend; b) they do not know they have an FFV; or c) they do not desire to use high-blend, for non-economic reasons. Based on regional dispensing station coverage and a fraction of non-users who are assumed to be amenable to using high-blends, FFV owners leak over time from the non-user to occasional user category. Under conditions of price parity between high-blend and regular gasoline, occasional users are assumed to fill 20% of their fuel requirements using high-blend. Regular users, on the other hand, are assumed to fill 80% of their fuel requirements using high-blend under conditions of price parity. Movement between occasional and regular users is driven by a long-term retail price differential between the two products.

The distribution of high-blend users provides a physical basis for ethanol usage among FFVs. Logit functions are used to translate relative high-blend/gasoline retail prices into instantaneous usage shares for both occasional and regular high-blend users. The distribution of occasional and regular users is then applied to these usage shares. The resultant user-weighted usage shares are

26

multiplied against potential high-blend fuel consumption in order to generate actual high-blend consumption within each region. Figure 24, below, provides a view into the logic involved.
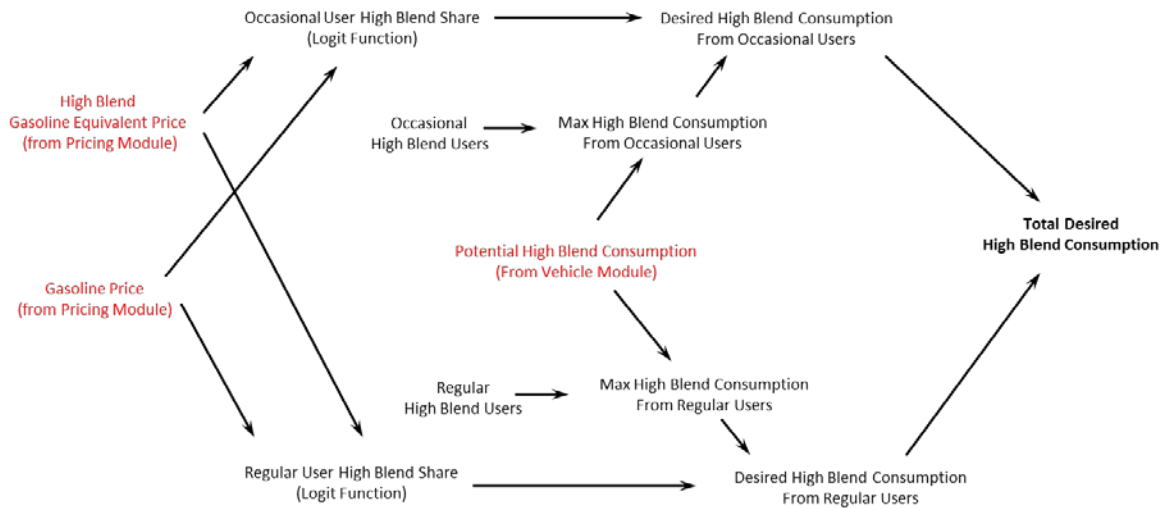
Figure 24. Logic behind high-blend consumption

Pricing and Inventory Module

The final module within the downstream sector accounts for ethanol pricing and inventory dynamics. Pricing and inventory for butanol, which in the model forms a substitute for ethanol in the low-blend market, are also captured within the pricing and inventory module.

In the model, ethanol inventory is aggregated across the entire supply chain within each region. The model allows for cross-regional movement of ethanol based upon regional surpluses or shortfalls within each region. The logic of inventory dynamics is shown in Figure 25.

There are several important features to note in Figure 25. First, note the three sources of regional ethanol production: the starch to ethanol module, the cellulose to ethanol module, and the import module. Second, note the regional import/export structure that facilitates cross-regional movement of ethanol. Third, note the single driver of ethanol consumption. This total reflects ethanol demand from both low-blend (i.e., E10) and high-blend (i.e., E85) uses. Finally, note the rich feedback that drives cross-regional movement of ethanol.

27
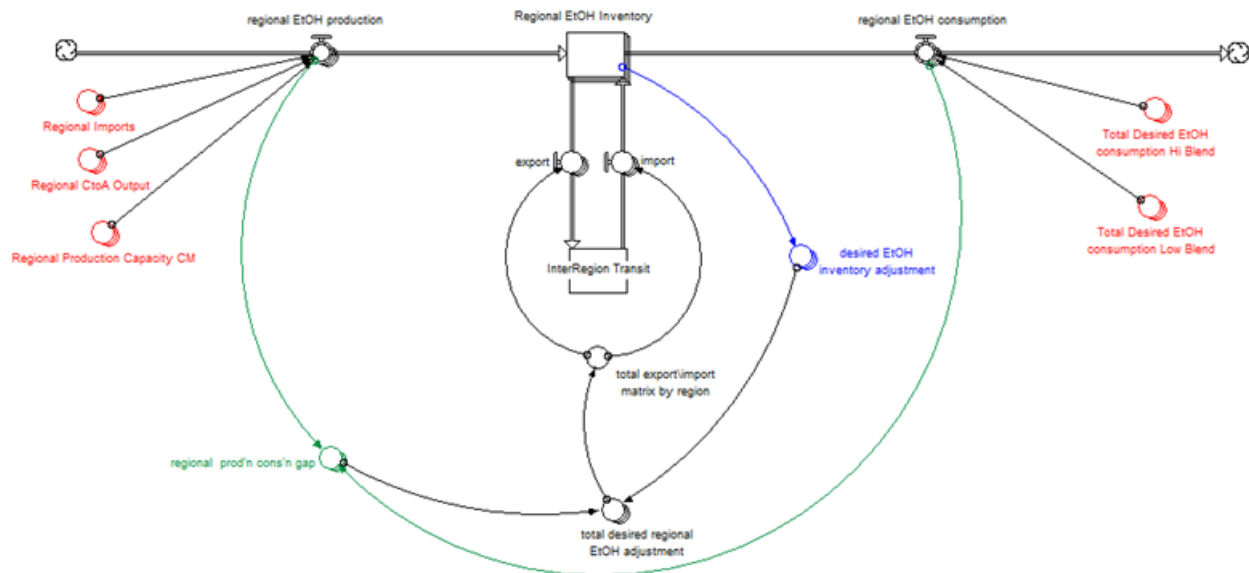
Figure 25. Downstream ethanol inventory dynamics

This cross-regional movement algorithm is relatively straightforward, and is outlined below:

- Calculate desired inventory adjustment in each region required to bring inventories to desired levels (blue connections in Figure 25)

- Calculate the regional production/consumption gap as the difference between regional production and consumption (green connections in Figure 25)

- Sum the inventory adjustment and production/consumption gap to arrive at overall desired movement in ethanol by region

- Roll up total desired imports and exports across all regions

- Limit total inter-regional movement to minimum of total desired imports, exports

- Allocate exports, imports in proportion to relative desired imports, exports.

Pricing for ethanol is considered at multiple downstream points along the supply chain. Figure 26 provides an overview of the approach. Ethanol price is calculated at point of production, at point of distribution, and at the pump. Supply/demand imbalances in the downstream supply chain drive changes in price at point of production (note that details around the pricing algorithm used within BSM are provided in Appendix A). Transport and storage costs, which vary based on distribution infrastructure within a region, are applied to the point of production price in order to generate an ethanol point of distribution price. The price for high-blend ethanol at the pump is determined as a weighted average of point of distribution price and gasoline prices, based on a regression analysis of the two. Not shown in Figure 27, but relevant to policy analysis, are multiple points along the supply chain where initiatives can work to reduce costs and/or change price as perceived by producers, distributors, retailers, or end users of ethanol or high-blend.

28
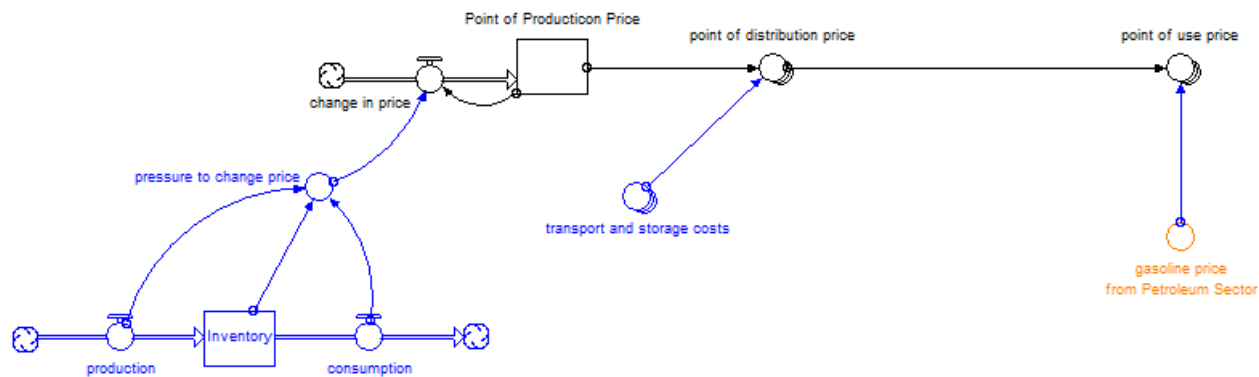
Figure 26. Simplified ethanol pricing structure

The pricing and inventory for butanol follows similar logic to that of ethanol, with some notable exceptions:

- A single, national inventory is considered

- In addition to its use in the low-blend gasoline market, butanol can be consumed for industrial uses

- Pricing for butanol is captured at point of production only. There is neither a point of distribution nor a point of use price for butanol.

The dynamics of butanol use and pricing center on the substitution of bio-butanol (produced within the BSM cellulose to butanol module) for butanol produced by other means, and on the substitution of butanol for ethanol in low-blend uses. These substitution dynamics are determined by relative price considerations. To capture these two dynamics, logit formulations are employed that translate relative prices into market shares. For industrial uses, the price of bio-butanol competes against an assumed alternative price of $4.00/gallon (this value can be varied as a scenario). For completion against ethanol, the endogenously-generated bio-butanol price is compared against the price of ethanol.

3.4 Vehicle Module

The primary purpose of the Vehicle Module in BSM is to provide inputs that represent potential demand streams for ethanol and for gasoline, from "regular" vehicles and from FFVs. These inputs are listed below.

- Regional potential low-blend consumption from FFVs

- Regional potential low-blend consumption from non-FFVs

- Regional potential high-blend consumption from FFVs

In order to provide these inputs to the rest of the model, we have developed a highly simplified accounting structure for vehicles of multiple types. Focusing on light-duty vehicles, this vintaging chain captures the cumulative impact of multiple scenarios around volume of new vehicles each year, new vehicle mix, new vehicle efficiency, vehicle miles traveled (VMT), and vehicle mortality. The model aggregates vehicles nationally. Regional population distributions

29

are used to apportion fuel consumption among the 10 USDA regions used by the model. In its operation, the module applies age-specific survivorship estimates to vehicles as they vintage through the chain. The model focuses on two distinct vehicle types within the light duty fleet:

1. Automobiles
2. Light trucks.

Within each vehicle type, ten technologies are considered:

    A. Gasoline

    B. Diesel

    C. Plug hybrid

    D. Hydrogen

    E. CNG

    F. FFV

    G. Gas HEV

    H. Gas PHEV

    I. Bi-fuel

    J. Other

For each of these 20 combinations, a scenario over time for new vehicle sales is accompanied by a scenario for new vehicle efficiency. The model tracks the implications over time of these new vehicle scenarios for overall vehicle efficiency and resultant fuel demand. The core model structure is shown below in Figure 27.

Each stage in the stock-flow chain represents a cohort of vehicles. Mortality flows remove vehicles from the system; vehicles that survive to the end of the cohort's time horizon are moved to the next cohort in the sequence. Cohorts 1-4 are each four years in duration. Cohort 5 contains vehicles that are 16 or more years of age. See Appendix C for details on the BSM approach to aggregating vehicles into four-year-sized lumps.

The parallel pathway, shown in Figure 27, accounts for the efficiency of vehicles in each cohort. Cohort-specific values for vehicles, efficiency, and vehicle miles traveled are used to calculate cohort-specific potential fuel usage, which is then summed over all cohorts to calculate overall potential fuel use.
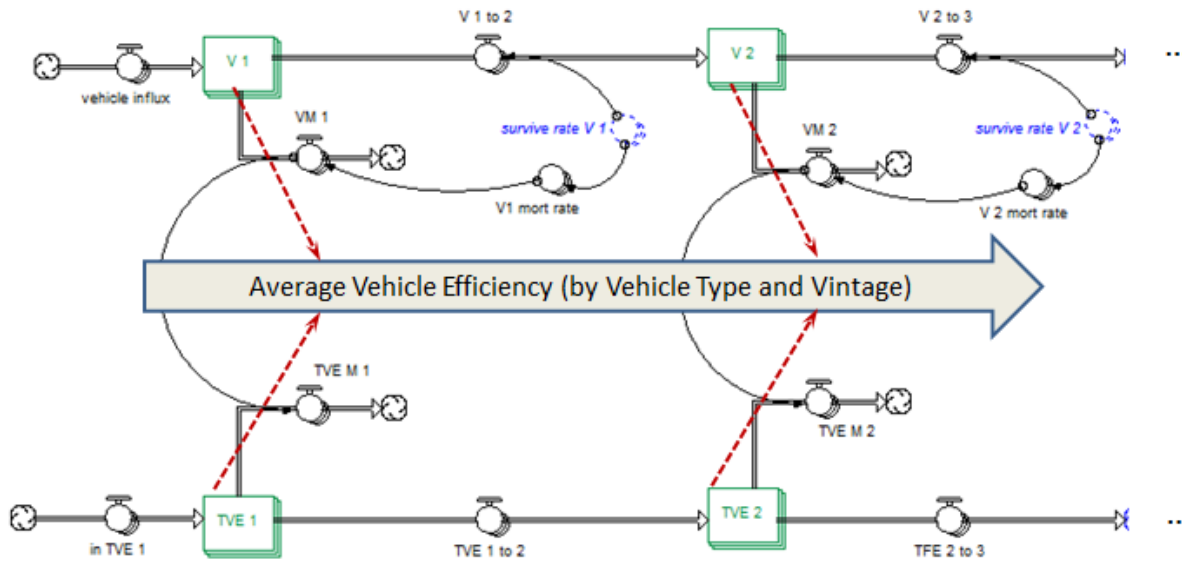
Figure 27. Tracking vehicles and efficiency (2 cohorts)

The Vehicle Module is designed to facilitate exploration of the cumulative impact resulting from changes in volume, mix, mortality, VMT, and efficiency. Structure in the model captures the *effects* of changes in fuel prices, consumer attitudes, and the like. While we have not provided an explicit representation of consumer choice mechanisms, in the Vehicle Module we have created the potential to develop internally consistent scenario sets in which vehicle inputs maintain a logical consistency with petroleum price scenarios. On the vehicle influx side levers within the Vehicle Module enable use of (or departure from) Annual Energy Outlook (AEO) projections for inflow volume, inflow mix, and efficiency (EIA 2014). Similarly it is possible to use or depart from National Energy Modeling System (NEMS) mortality rates, and to use or modify AEO VMT scenarios (EIA 2009).

3.5 Oil Industry Sector

The oil industry sector is relatively simple, containing a single module that houses:

1. A set of scenarios used to determine crude oil prices

2. Refinery product prices for diesel, jet fuel, and gasoline

3. Algebraic relationships that translate crude oil prices, refinery product prices, and an assumed refinery "drop in point" for each fungible fuel pathway into price inputs for the different conversion modules and for the downstream pricing and inventory module

4. Accounting structure that captures petroleum displaced by diesel, jet fuel, and gasoline produced by the different fungible fuel pathways.

3.6 Import Module

The Import Module is an exceedingly simple structure focused on the import of fuel ethanol from outside U.S. borders based on relative price considerations. Figure 28 below shows the essential structure of the module.

This structure compares the ethanol point of production price generated within the downstream pricing and inventory module against a threshold (including tariffs) that reflects the cost of bringing fuel ethanol into the United States. As the price within the United States exceeds the threshold, and increasing fraction of offshore production capacity is utilized. This simple structure enables analysis of scenarios around tariff policies, cost reduction, and capacity growth for offshore ethanol production facilities.
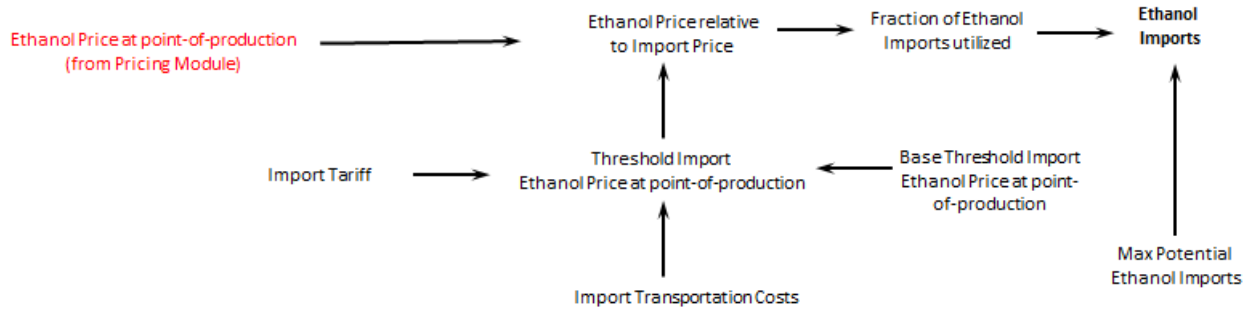
Figure 28. Logic behind ethanol imports calculation

4 Interconnections Among Sectors

Figure 2 provided a high-level overview of the sectors that comprise BSM, and the previous discussion has given a detailed view into the modules that are found within each sector. Another perspective on the system is given by the nature of the interconnections among the different sectors. As shown in Table 3, the connections between sectors are relatively few in number, typically consisting of price signals and supply/demand quantities.

Table 3. Inter-sector Connections

From/To	Feedstock Supply & Logistics	Conversion	Import	Oil Industry	Downstream
Feedstock Supply & Logistics		Feedstock consumption Feedstock price (plant gate)			
Conversion	Feedstock demand Cost to price ratios Output capacity			Infrastructure-compatible fuel production by pathway	Ethanol Production Butanol Production
Import					Ethanol import
Oil Industry	Gasoline point of distribution price	Module-specific price input			Gasoline point of distribution price
Downstream		Ethanol point of production price Butanol point of production Price input	Ethanol price input		
Vehicles					Potential lo-blend cons'n from FFV Potential lo blend cons'n from non-FFV Potential hi-blend cons'n Potential gasoline cons'n

33

5 Data Inputs

Multiple data inputs are required to run BSM. These inputs range from agricultural cost and yield parameters, to performance and learning parameters for the various conversion modules, to logit coefficients, to petroleum prices, to adoption rates for new farm practices and for dispensing station owners. Given the forward-looking nature of BSM, it is not surprising that the availability and quality of input data is highly variable. In many instances, assumptions or informed opinion were used to populate the parameter space. Table 4 provides a summary of the data inputs used in developing BSM.

Table 4. Summary of Data Inputs to BSM

	Input data area	Source(s)	Comments
Feedstock Supply & Logistics	Crop production costs	ORNL/POLYSYS	Assumed constant over simulation time frame Ongoing interaction with ORNL analysts
	Energy price/crop production price coupling	Pacey study (McNulty 2010)	Price coupling factors derived from econometric analysis
	Yields	ORNL/POLYSYS	Yield growth treated as assumption/scenario
	Calibration data for production, prices	USDA baseline (United States Department of Agriculture n.d.)	Calibration done ~annually based on annual updates to baseline and updates to input data from ORNL
	Logit parameters	Assumption	Assumptions modified as needed as part of calibration process
	Harvest, transportation, Q&H, preprocessing logistics	INL/Biomass Logistics Model (BLM)	Structure and input data updated periodically to reflect ongoing interaction with INL analysts
Conversion	Performance and cost data for different conversion options	NREL design reports PNL design reports Analysis papers Expert opinion Internal secondary analysis/interpolation	NREL staff are assembling and vetting these data. For some conversion options, formal analysis reports do not exist. We are in process of vetting available data, developing assumptions, and facilitating an expert review
	Learning curve parameters	Assumption informed by Beck study (Beck 2010)	Sensitivity analysis planned
	Logit parameters	Assumption	Plan sensitivity and robustness analysis around current logit parameters
	Construction capacity	Assumption	
	NPV of "other" option	Assumption	
Oil Industry	Oil price	EIA (United States Energy Information Agency n.d.), arbitrary scenarios	Data taken from "official" scenarios. Oil price shocks, other scenarios available to the system
	Fuel mix	Assumption	Treated as assumption but informed by design reports
	Drop in points	Assumption	
Downstream EtOH	Distribution Terminals by region	EIA	Developed from EIA data in 2008
	Initial mix of terminals with/without infrastructure	Assumption	
	Infrastructure acquisition rate	Assumption	
	Number, distribution of dispensing stations by ownership, dispensing station economics	NREL (Johnson and Melendez 2007 draft) NACS (National Association of Convenience Stores 2007)	NACS provides a rich perspective on "other" sales associated with dispensing stations in the spreadsheets that accompany the text of their annual report
	Initial repurposable stations	Assumption	
	Station adoption rates as f(NPV)	Assumption	

	Input data area	Source(s)	Comments
	Logit parameters for fuel use	assumption	
	Vehicle influx, miles traveled, miles per gallon	EIA/NEMS	
	Ethanol price at point of distribution	Assumed	Assumed values for storage and transport applied to endogenous point of production price
	High blend point of use price	NREL/Lexidyne regression	Regression of available data provides weighting factors for point of use price
Import	Capacity, price threshold for import, learning curve parameter	Assumed	Values used to calibrate against observed data for fuel ethanol imports

6 Analysis

Analysis efforts with BSM have been ongoing since early 2010. Initial analysis efforts used the earlier version 2.0 of BSM, a model of the biofuels supply chain that focuses on ethanol from cellulose and starch crops. Beginning in the fall of 2011, the BSM team has been designing and conducting analyses using the October 2012 version of the model, which focuses on both ethanol and infrastructure-compatible fuels. Past, current, and prospective analysis efforts include the following:

- Team vetting of model structure and parameters
    - Equation-by-equation review of model structure
    - Automated factorial sensitivity and robustness analysis
- Revision of model structure and parameters as required to address issues uncovered in team vetting process
- Calibration of model using 2012 input and calibration data sets
    - Agricultural system input data
    - Feedstock logistics data
    - Cost and performance data for conversion options
    - Fuel price scenario inputs
- Sensitivity studies around
    - Nth plant capital cost
    - Learning parameters—initial conditions and progress ratios for different conversion options/development stages
    - Initial vs. Nth plant multipliers for capital cost and yield
    - Pilot and demo scale scenarios for different conversion options
- "Generic" policy studies
    - Impact of capital cost reduction policies
    - Impact of subsidies aimed at various stages of the supply chain
- Targeted policy studies—aimed at providing a perspective on a variety of initiatives under consideration by DOE and other organizations.

The BSM project utilizes modern software-engineering methodologies to maintain model quality and enable analysis flexibility. An open-source configuration management and version control system, named Subversion, is used to track changes in the BSM model, documentation, and other project-related files. Documentation and metadata for variables are embedded directly in the STELLA code. Raw data sources are archived, provenance/pedigree metadata is tracked, and input data are stored (and ultimately processed) within a relational database. Multidimensional data analysis, statistics, and visualization tools are linked to the database in an architecture that allows for the automated "refresh" of visualizations and analyses when new scenarios are run.

This computing infrastructure supports a high-throughput analysis process such as a "design-of-experiments" approach for simulation studies that involve complex combinations of policy scenarios, sensitivity analysis, and uncertainty quantification. The automation of simulation studies involves retrieving input parameters from the database, running STELLA models in "batch mode", and then storing simulation outputs into the database. Furthermore, the BSM source files can be copied and run on multiple machines simultaneously. This centralized system allows any BSM team member to re-create any past scenario—by viewing previous runs, finding the corresponding model on the model repository, or even re-running previous versions of the BSM with updated input parameters. Simulation outputs can then be imported into any graphics software to visualize results, analyze trends and develop insights.

An important part of the BSM analysis plan is the dissemination of results. Accordingly, at this writing, multiple papers and articles are under development, submitted for review, or published. Additionally, analyses reports have been archived in the BioEnergy Knowledge Discovery Framework (https://www.bioenergykdf.net).

7 Concluding Remarks

The Biomass Scenario Model provides a rich representation of the supply chain associated with the production of biofuels. By integrating feedstock production and logistics, multiple conversion options, and market dynamics for butanol, fuel ethanol, and fungible fuels (gasoline, diesel, jet fuel), the model serves as a vehicle for exploring the mechanisms by which the biofuels industry might develop beyond its current state. By providing a representation that reflects both the physics and economics of the system, BSM can serve as a tool for building understanding around initiatives that seek to stimulate sustainable development of the industry. By representing the system of interactions simply and transparently, the model can shed light on gaps in the data as well as areas where understanding of system structure is in need of enrichment.

Analyses of earlier versions of the model—both as standalone modules and in integrated form—have underwritten powerful insights about the nature of the supply chain and of the nature of policy initiatives required to stimulate industry take off. We anticipate that the October 2012 version of the model will facilitate the development of additional insight and understanding around the evolution of the biofuels industry.

References

Beck, R.W. (2010). *Biorefinery Learning Curve Analysis*. Golden, CO: National Renewable Energy Laboratory.

Hax, A.C; Majluf, N.S. (1982). "Competitive Cost Dynamics: The Experience Curve." Interfaces (12:5); pp. 50-61.

Johnson, C; Melendez, M. (2007). *E85 Retail Business Case: When and Why to Sell E85*. NREL/TP-540-41590. Golden, CO: National Renewable Energy Laboratory.

McNulty, M.S. (2010). *Energy Price-Biofuel Production Cost Coupling Analysis*. AEU-9-88344-01. Boulder, CO: Pacey Economics Group.

Merrow, E.W. (1983). *Cost Growth in New Process Facilities*. P-6869. Santa Monica, CA: Rand Corporation.

National Association of Convenience Stores. (2007). *Annual Report*. Alexandria, VA: National Association of Convenience Stores.

Newes, E.; Inman, D.; Bush, B. (2011). "Understanding the Developing Cellulosic Biofuels Industry through Dynamic Modeling," Chapter 18. Dos Santos Bernardes, M.A., ed. Economic Effects of Biofuel Production. InTech Open Access.

Peterson, Steve et al. (2013). *An Overview of the Biomass Scenario Model.* 31st International Conference of the System Dynamics Society. Cambridge, Massachusetts.
Accessed August 2015: http://www.nrel.gov/docs/fy15osti/60172.pdf

Richmond, Barry. 2004. isee systems, inc. iseesystems.com

Sterman, J.D. (2000). *Business Dynamics: Systems Thinking and Modeling for a Complex World*. New York: Irwin/McGraw-Hill.

Train, K.E. *Discrete Choice Methods with Simulation*. (2003). Cambridge, U.K.: Cambridge University Press.

U.S. Energy Information Administration (EIA). (2009). *The National Energy Modeling System: An Overview*. DOE/EIA-0581(2009). Washington, D.C.: Energy Information Administration. Accessed July 2014: http://www.eia.gov/oiaf/aeo/overview/.

EIA. (2014). *Annual Energy Outlook 2014 with projections to 2040.* Washington, D.C.: Energy Information Administration. Accessed July 2014: http://www.eia.gov/forecasts/aeo/.

Vimmerstedt, L., B. Bush, and S. Peterson. "Ethanol Distribution, Dispensing, and Use: Analysis of a Portion of the Biomass-to-Biofuels Supply Chain Using System Dynamics." PLoS ONE (7:5); e35082.

Westcott, P.; Trostle, R. (2014). *USDA Agricultural Projections to 2023*. USDA Agricultural Projections No. (OCE-141). Washington, D.C.: U.S. Department of Agriculture. Accessed July 2014.

Appendix A. Pricing Within BSM

In BSM, an endogenous pricing mechanism is an essential component of the structure that underwrites industry development. The model incorporates endogenous pricing structures for

- Each of the commodity crops (corn, wheat, cotton, small grains, soy)
- Hay (regional markets)
- Cellulosic feedstocks (regional markets)
- Ethanol
- Butanol

It can be helpful to view price mechanisms within BSM as central components of an economic control system. Each price signal evolves in response to the interplay of the forces of supply and demand. As production, consumption, and inventories change over time, price responds to imbalances. Prices, in turn, play a critical role in the investment, allocation and utilization decisions of producers of agricultural products and of biofuels. They also play a critical role in the fuel use decisions for butanol and for high-ethanol-blend fuels.

In developing the pricing structure used in BSM, we were mindful of multiple design constraints. First, the pricing mechanism needed to be simple so as to be understandable to a broad audience of model users. Second, the structure needed to be sophisticated, in order to *not* generate spurious dynamics. A simplistic pricing formulation can lead to steady-state error in controlled quantities or can become trapped in unrealistic states in response to extreme condition tests. Finally, the pricing mechanism needed to be flexible enough to support real-world circumstances such as market initiation and scale-up.

The basic feedback relationships of the BSM pricing mechanism are shown in Figure 29.
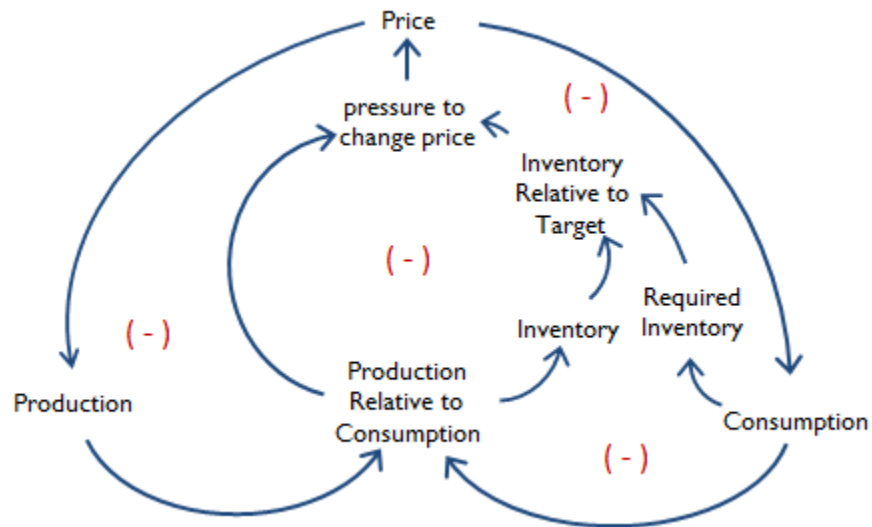
Figure 29. Stylized view of feedbacks in BSM pricing mechanism. (-) indicates negative feedback loop

In this simple diagram, price works to balance production and consumption *and* to balance inventory against desired or target levels. Production/consumption imbalances create pressure to change price, as do imbalances between inventory and target inventory (which, in turn, depends on consumption). In order to accumulate or integrate pressure over time, price must be represented as a stock. The representation of price as a stock, in conjunction with pressure from inventory, results in oscillatory tendencies in the system; oscillations are dampened by the presence of feedback connections around production, consumption, and price.

Figure 30 shows output from a simplified model of pricing/inventory/producer/consumer dynamics, which uses the basic pricing structure found in BSM. The test shows the equilibrium-seeking tendencies of the structure, in response to a 10% shift in product demand.
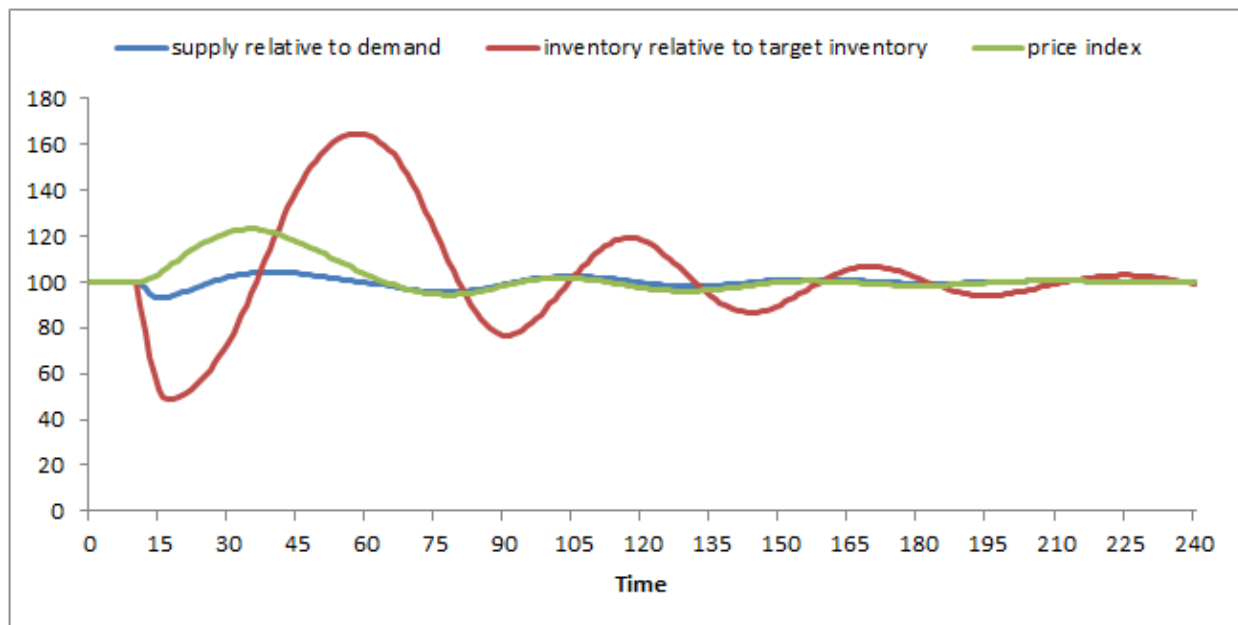
Figure 30. Response of pricing system to 10% step-increase in demand

In BSM pricing, the mechanisms that connect production, consumption, and inventory to fractional change in price are significantly more detailed, as illustrated in Figure 31. The structural arrangement shown in Figure 31 is used for throughout BSM, including:

- Ethanol at point of production
- Commodity crops
- Hay
- Cellulosic feedstocks
- Butanol.

43
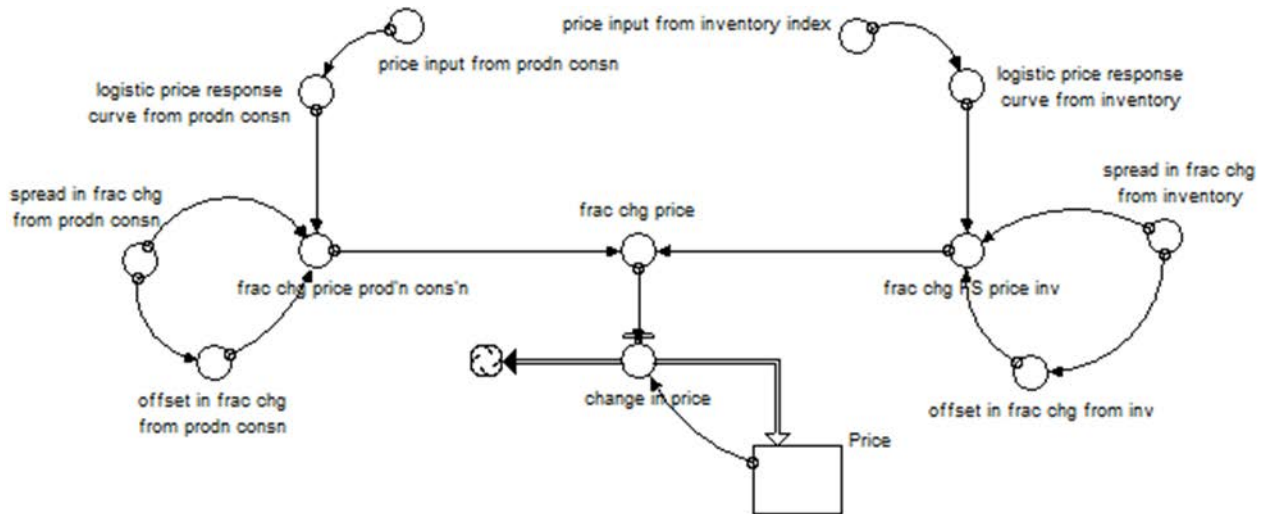
Figure 31. Detail of BSM generic pricing structure

The algorithm associated with this structure uses a bit of sophisticated math, but is relatively straightforward. It begins by calculating the price input—either from inventory or from production relative to consumption—as a distance from equilibrium in doublings. When the ratio is 1, the input is at its equilibrium value. When it is 2, it is one doubling away from equilibrium. When it is 0.5, it is one halving away from equilibrium. To capture this distance simply, the model uses logarithm functions as illustrated in Figure 31.

Second, the price input processed through a logistics function to generate a well-behaved response curve. Price input and logistics calculations are shown in Figure 32.
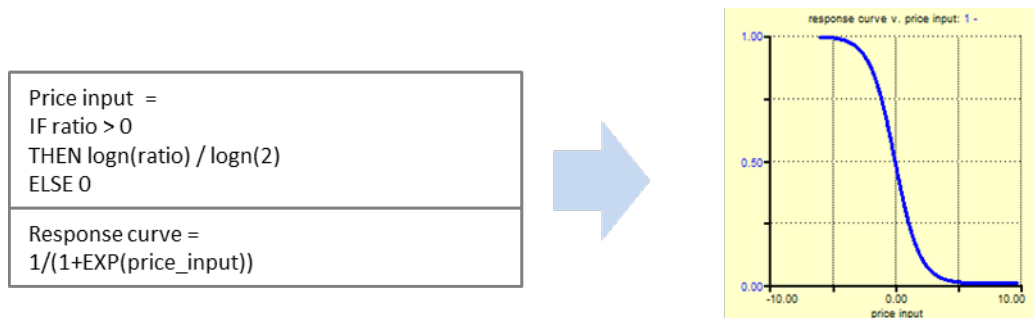
Figure 32. Illustrative price input and response curve calculations

The third step in this algorithm is to scale the response curve by shifting its intercept to (0,0) and setting its asymptotes to desired maximum and minimum fractional changes in price. Finally, the total fractional change in price is calculated as the sum of fractional changes from inventory and production/consumption, and the result is applied to the price to generate a total fractional change in price.

In BSM, this generic pricing structure is applied to multiple market situations. Context-specific details, outlined below, are important to note.

44

Ethanol Pricing:

- Price index is applied to baseline price to generate point of production price.

- Inventory target is calculated as a multiple of desired ethanol consumption—enabling inventory to grow as demand for ethanol grows.

- Price feedback affects consumption of ethanol through:

    o Buildout of dispensing station high-blend tankage and equipment

    o Relative price dynamics with gasoline in high-blend market for occasional and regular users

    o Relative price dynamics with butanol in low-blend market

- Price feedback affects production of ethanol through:

    o Investment attractiveness

    o Utilization of existing conversion facilities

    o Ethanol import dynamics

Butanol Pricing:

- Price index is applied to baseline price to generate point of production price.

- Inventory target is calculated as a multiple of desired butanol consumption—enabling inventory to grow as demand for butanol grows.

- Price feedback affects consumption of ethanol through:

    o Relative price dynamics with non-bio-butanol in the industrial market

    o Relative price dynamics with ethanol in low-blend market

- Price feedback affects production of butanol through:

    o Investment attractiveness

    o Utilization of existing conversion facilities

Cellulosic Feedstocks:

- A single aggregate average feedstock price is considered at the plant gate in each of 10 regions.

- Inventory targets are determined as a multiple of desired feedstock consumption, enabling inventory to grow as demand grows.

- Grower payment considers plant gate price net of harvest and logistics costs.

- Feedstock price is set to zero until a conversion facility enters the development pipeline. Assumed "seed" values for feedstock prices are used to set the initial nonzero price.

- A price ceiling for cellulosic prices is set by downstream conversion facility dynamics. This ceiling constrains prices from growing in excess of their value in producing end products.

45

- Price feedback affects consumption of feedstocks through:
    - Investment attractiveness for conversion facility investment
    - Utilization of existing conversion facilities
- Price feedback affects production of feedstocks through:
    - Movement along supply curve for forest, urban residues
    - Relative grower payment for producers of feedstocks, which in turn changes land allocation
    - Changing the rate at which producers become producers of cellulose

Commodity crops:
- A single national aggregate average price is used for each commodity crop.
- An adjustment factor is used to couple small grain prices to corn. These product classes are substitutes for one another in the animal feed market, and their prices tend to move in concert.
- Inventory feedback is developed using stock:use ratios. These ratios are allowed to float over time, enabling the model to capture observed history around prices and stock:use ratios.
- Price feedback affects consumption of commodity crops through small changes to baseline USDA projections.
- Price feedback affects production of commodity crops through relative grower payments to producers, which in turn changes land allocation.

Hay:
- Regional price indices are used to represent hay .
- Inventory dynamics are not included as an input to hay prices, reflecting the assumption of "use it or lose it" for hay.
- Demand and supply for pasture are considered in hay price dynamics.
- Price feedback affects consumption of hay through small changes to baseline model values.
- Price feedback affects production of hay and pasture through relative grower payments to producers, which in turn changes land allocation.

46

Appendix B. Logit as Allocation Mechanism Within BSM

In BSM, logit functions are a mechanism for allocating resources among multiple competing uses. Detailed discussion of the logit function can be found in a variety of texts and articles dealing with consumer choice. For example, Train (Train 2003) provides a thorough introduction to the logit, generalized extreme value, and a wide range of other approaches. The logit function expresses the likelihood P of choosing alternative i from the set of J alternatives given an observed utility of x. A simple form of the logit is shown below:

$$P_i = \frac{e^{(k_i + Bx_i)}}{\sum_j (e^{(k_j + Bx_j)})}$$

The parameter k reflects unobserved or unexplained utility, while the parameter B is a scaling factor. The logit function has several desirable characteristics. Among them:

- The function can be interpreted in terms of the utility associated with alternatives within a set of choice.

- The sum of probabilities across all choices is 1.

- There is a sigmoid relationship between utility and the resultant probability, which is beneficial under extreme conditions.

The typical interpretation of the logit formulation, in the context of consumer choice, is the probability of choosing a particular alternative. In BSM, this probabilistic interpretation is applied to a *population* of actors (for example, farmers, investors in conversion facilities, consumers as they are deciding to fuel their vehicles) in order to generate an aggregate allocation of land use, investment, or fuel use.

Logit formulations can be found throughout BSM. Table 5 summarizes these uses.

Table 5. Uses of Logit Formulation Throughout BSM

Module	Usage	Dynamic Inputs	Notes
Feedstock Supply	Crop land allocation • Commodity crops ○ With/without residues • Perennial cellulosic crop • Hay Pasture land allocation • As pasture • As pasture harvested as cellulosic feedstock • Perennial cellulosic energy crop	Per-acre grower payment for respective uses.	For crop land, nested logit function is used to allocate among broad groups (e.g., commodity vs. perennial cellulosic vs. hay) and then among different commodity crops
Conversion and Relative Attractiveness	Allocation of facility construction resources among alternate pioneer and commercial scale conversion pathways in different regions	NPV of respective conversion pathways	Nested logit function is used to allocate construction capacity among different conversion platforms (e.g., fast pyrolysis) and then among different regions
Pricing and Inventory (Downstream)	Displacement of EtOH by butanol in low-blend mixes Displacement of non-bio-butanol in industrial market	Butanol, EtOH prices	
Fuel Use	Allocation of fuel sales between high-blend and gasoline	Price of gasoline Price of high-blend	

Appendix C. Aggregation of age Classes in the Vehicle Module

The October 2012 version of the Vehicle Module, like the other modules within BSM, reflects design tradeoffs between the competing pressures of detail "realism" and usability. It is conceptually straightforward to create a model containing great detail around vehicle type, regional distribution, and age distribution of vehicles. Unfortunately, the computational overhead required to simulate this detail would quickly become unmanageable. In an earlier version of BSM which incorporated this detail, we were required to run the Vehicle Module separately from the rest of the model and then import fuel demand scenarios separately.

In this version of the model, we have reduced computational overhead significantly by aggregating age distribution of vehicles. We represent vehicle vintages using five distinct cohorts. Each cohort represents four years of vehicle life. Within each cohort, each year, vehicles are scrapped or they get older. Those vehicles that survive to the end of a cohort are transferred to the next cohort. A portion of the vehicle aging logic is shown in Figure 33.
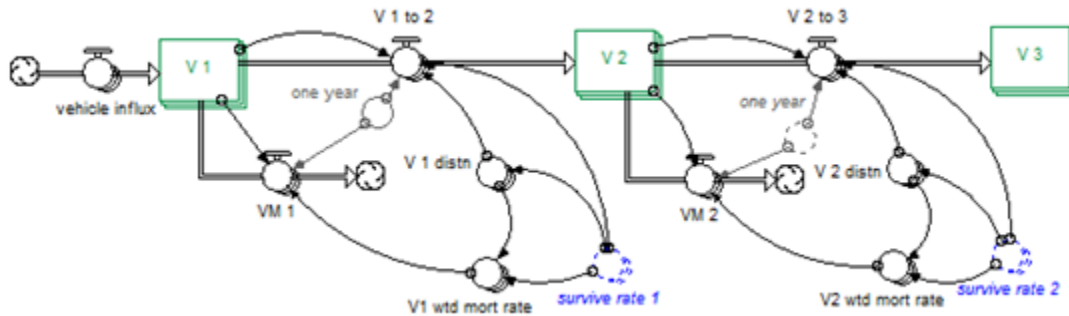

Figure 33. Structure of vehicle vintaging

This structure aggregates together vehicles of multiple ages, and it is important to provide a reasonable estimate of the distribution of vehicles within each cohort. To do so, we consider the age-specific survival rates within each cohort, using these to derive an approximation of distribution of vehicles across the cohort:

Let S_n = survival rate for year n in cohort, $0 <= S_n <= 1$, $S_0 = 1$

D_n = fraction of cohort population in year n

$D_1 = S_0 / (S_0 + S_0 * S_1 + S_0 * S_1 * S_2 + S_0 * S_1 * S_2 * S_3)$

$D_2 = S_0 * S_1 / (S_0 + S_0 * S_1 + S_0 * S_1 * S_2 + S_0 * S_1 * S_2 * S_3)$

$$D_n = \frac{\prod_{i=0}^{n-1} S_i}{\sum_{i=0}^{n-1} \prod_{i=0}^{n-1} S_i}$$

$D_3 = S_0 * S_1 * S_2 / (S_0 + S_0 * S_1 + S_0 * S_1 * S_2 + S_0 * S_1 * S_2 * S_3)$

$D_4 = S_0 * S_1 * S_2 * S_3 / (S_0 + S_0 * S_1 + S_0 * S_1 * S_2 + S_0 * S_1 * S_2 * S_3)$

Age-specific survival rates are then applied to this distribution of vehicles in order to calculate distribution-weighted age-specific mortality rates, which are then summed and applied to the number of vehicles in the cohort to generate a mortality flow. The survival rate for the last year

49

in the cohort is applied to the appropriate distribution in order to generate movement of vehicles to the next cohort.

Figure 34 compares the transient response of a single four-year cohort of the BSM vintaging structure against a simple one-stock structure, and against a more disaggregated structure with four one-year cohorts. For both systems, yearly survival rates are set to 50%. In the test, both systems are initialized at zero. Inflow to each system is set to 100 initially; the inflow steps down to 50 at time 10.
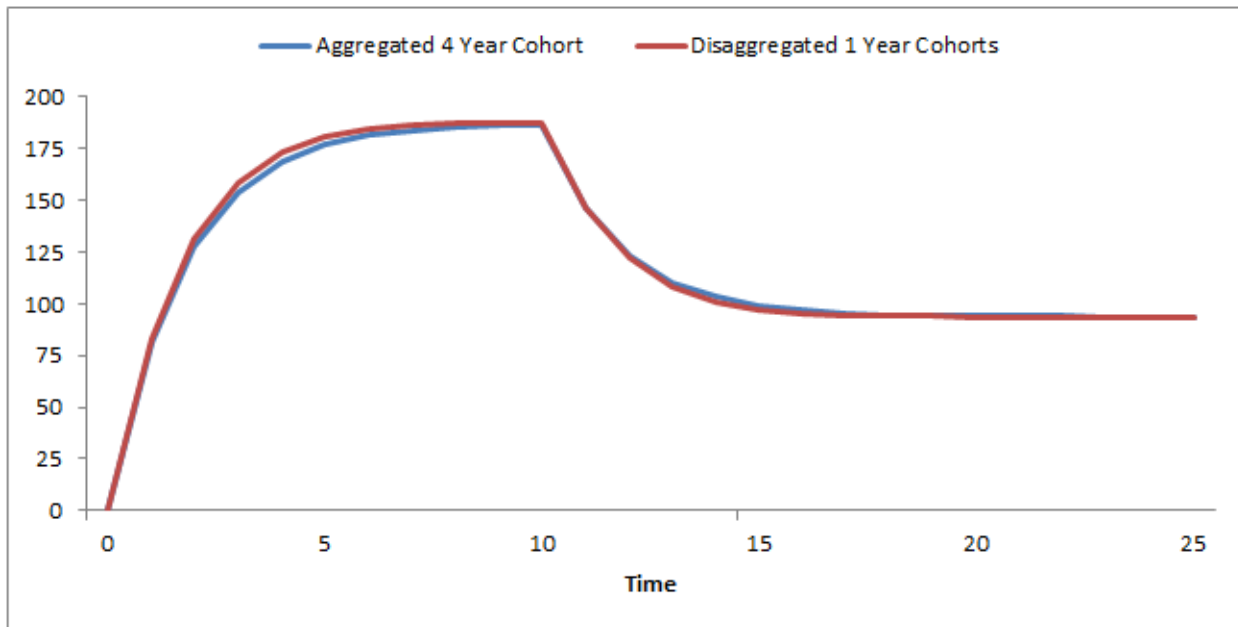
Figure 34. Comparison of BSM and disaggregated vehicle cohorts