An Overview of the Biomass Scenario Model

Steve Peterson, Lexidyne LLC

Emily Newes, NREL

Danny Inman, NREL

Laura Vimmerstedt, NREL

David Hsu, NREL

Corey Peck, Lexidyne LLC

Dana Stright, Lexidyne LLC

Brian Bush, NREL

August 2013

Submitted to

The 31st International Conference of the System Dynamics Society

Cambridge, Massachusetts USA

July 21 – July 25, 2013

Introduction

Biofuels are promoted in the United States through aggressive legislation, as one part of an overall strategy to lessen dependence on imported energy as well as to reduce the emissions of greenhouse gases (Office of the Biomass Program and Energy Efficiency and Renewable Energy, 2008). For example, the Energy Independence and Security Act of 2007 (EISA) mandates 36 billion gallons of renewable liquid transportation fuel in the U.S. marketplace by the year 2022 (U.S. Government, 2007). Meeting such large volumetric targets has prompted an unprecedented increase in funding for biofuels research, much of it focused on producing ethanol and other fuel types from cellulosic feedstocks[1] as well as additional biomass sources (such as oil seeds and algae feedstock). In order to help propel the biofuels industry, the U.S. government has enacted a variety of incentive programs (including subsidies, fixed capital investment grants, loan guarantees, vehicle choice credits, and aggressive corporate average fuel economy standards) -- the short- and long-term ramifications of which are not well understood. Efforts to better understand the impacts of incentive strategies can help policy makers to develop a policy suite which will foster industry development while reducing the financial risk associated with government support of the nascent biofuels industry.

Purpose and overview

This paper describes the Biomass Scenario Model (BSM), a system dynamics model developed under the support of the U.S. Department of Energy (DOE). The model is the result of a multi-year project at the National Renewable Energy Laboratory (NREL). It is a tool designed to better understand biofuels policy as it impacts the development of the supply chain for biofuels

[1] These feedstocks, such as agricultural and forestry residues, perennial grasses, woody crops, and municipal solid wastes, are advantageous because they do not necessarily compete directly with food, feed, and fiber production.

in the United States. In its current form, the model represents multiple pathways leading to the production of fuel ethanol as well as advanced biofuels such as biomass-based gasoline, diesel, jet fuel, and butanol).

The BSM uses a system dynamics modeling approach (Bush et al., 2008), developed using the STELLA software platform (isee systems, 2010) to model the entire biomass-to-biofuels supply chain. 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, as shown in Figure 1.

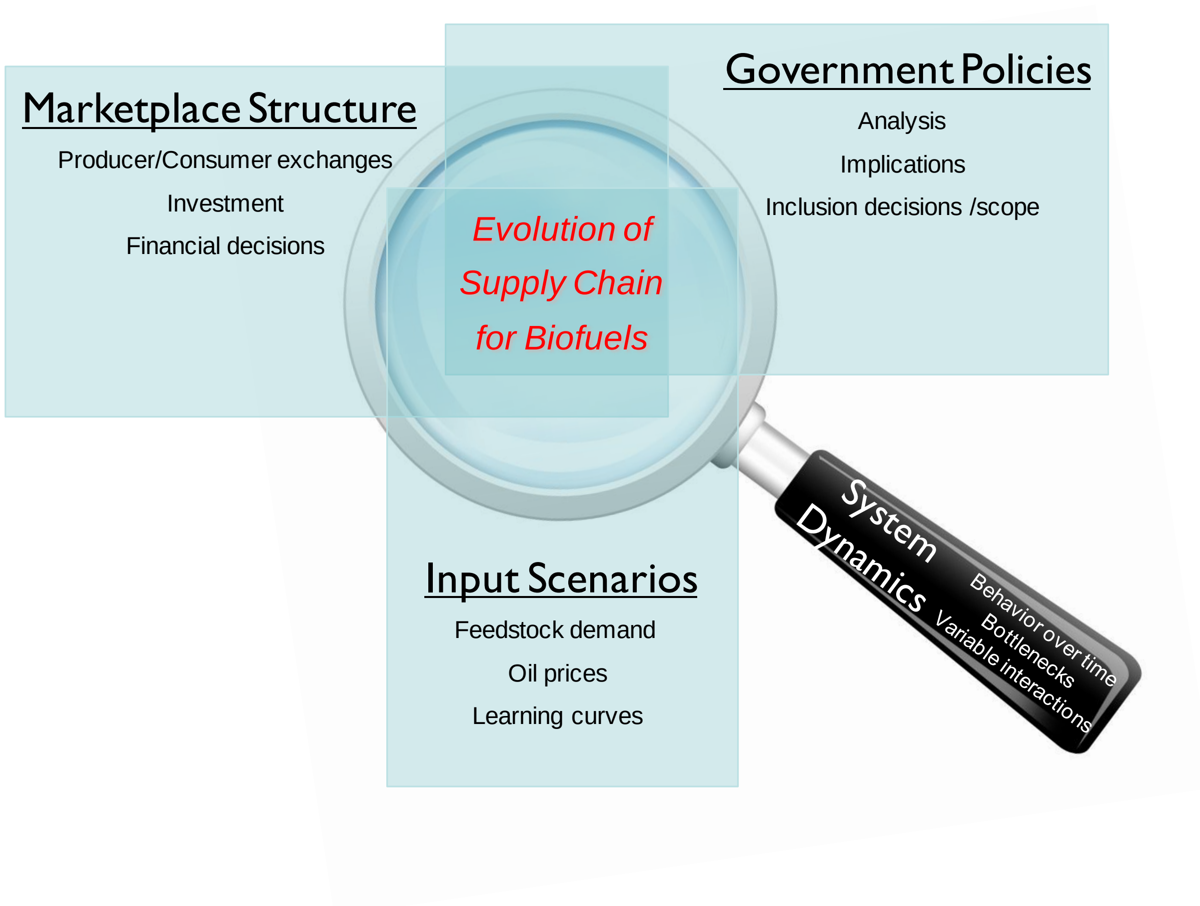

Figure 1. BSM strategy and approach

In this paper we begin with a description of the BSM architecture. Then, we provide a more detailed view of the sectors and modules which comprise the BSM. Third, we outline the "back-end" system that we have developed in order to support analysis efforts with the BSM. Fourth, we describe a set of scenarios used as the basis for policy exploration with the model. Finally, we provide a summary of current and potential uses for the model. A set of appendices detail important structures relating to pricing, investment, and vehicle vintaging.

An overview of the BSM architecture

BSM has been designed in a top-down, modular fashion which allows material (feedstocks) to flow down the supply chain and be converted into various types of biofuels, with feedback mechanisms among and between the various modules. In developing the model, we have taken care to create a structure that is transparent, modular, and extensible, enabling standalone analysis of individual model components as well as testing of different module combinations. As shown in Figure 2, the model is framed as a set of interconnected sectors and modules.

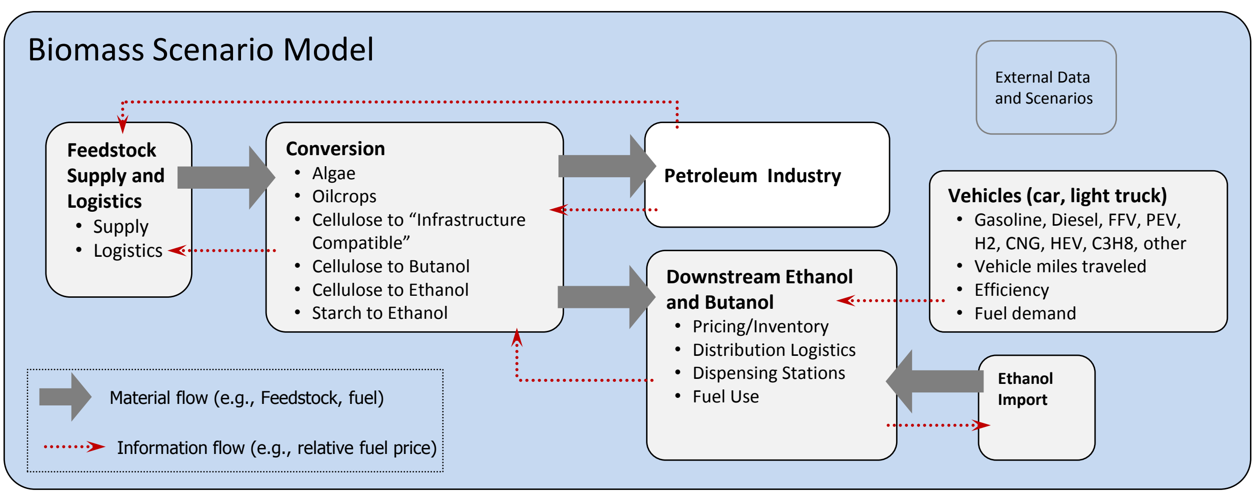

Figure 2. Overview of BSM structure

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 U.S. agricultural system. It incorporates harvesting and transportation logistics associated with cellulosic feedstock, as well as feedstock supply and logistics associated with forest, urban, and agricultural residues.

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

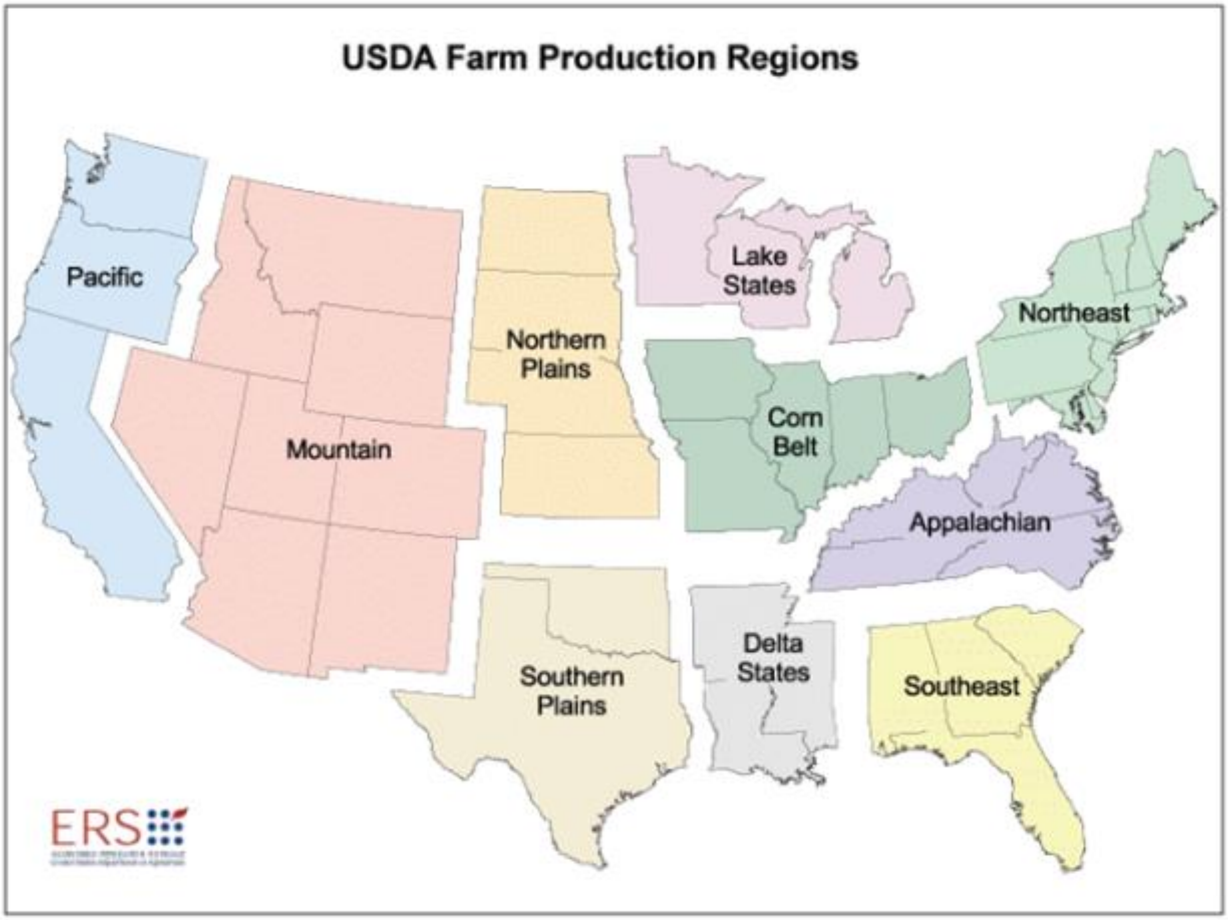

Figure 3. USDA farm production regions

Conversion

The conversion sector is composed of six different 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. The 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 the oxygenate market. 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 dynamics, investment decision logic, and utilization logic for both pioneer- and full-commercial scale facilities.
- Cellulose to “refinery ready:” This module captures the industry development of cellulose-to-refinery-ready “infrastructure compatible” conversion processes. The model structure can accommodate the following conversion options:
    - Fast pyrolysis
    - Fischer-Tropsch
    - Methanol to gasoline
    - Catalytic pyrolysis
    - Fermentation
    - Aqueous phase reforming.

As with other cellulosic conversion 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.

Petroleum industry

The petroleum industry sector comprises scenario inputs around crude oil prices, providing 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.

Downstream ethanol and butanol

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

- 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" oxygenate market and the "high-blend" market associated with flexible-fuel vehicles (FFV). Bio-butanol is assumed to serve as a substitute for ethanol in the oxygenate 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 flexible-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.

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.

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 abroad. 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.

A more detailed view of BSM[2]

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

[2] The following sections provide a "deep dive" into specifics related to each sector/module in the BSM. The casual reader may wish to move ahead to the "Interconnections between sectors" section.

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, as summarized in Table 1.

Table 1. Summary of feedstocks produced by feedstock supply module

Feedstock	Source	Use	Notes
Corn	Cropland	Ethanol	
Soy	Cropland	"Refinery-ready" fuels	
Other oil seed	Cropland (small grains)	"Refinery-ready" fuels	Model does not explicitly represent land allocation to other oil seed (e.g., rapeseed)
Crop residue	Cropland	Ethanol | butanol | "refinery-ready" fuels	Model allows residue collection from corn, wheat, other grains, cotton
Herbaceous cellulosic energy crop	Cropland pastureland	Ethanol | butanol "refinery-ready" fuels	
Woody cellulosic energy crop	Cropland pastureland	Ethanol | butanol "refinery-ready" fuels	
Pasture	Pastureland	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.

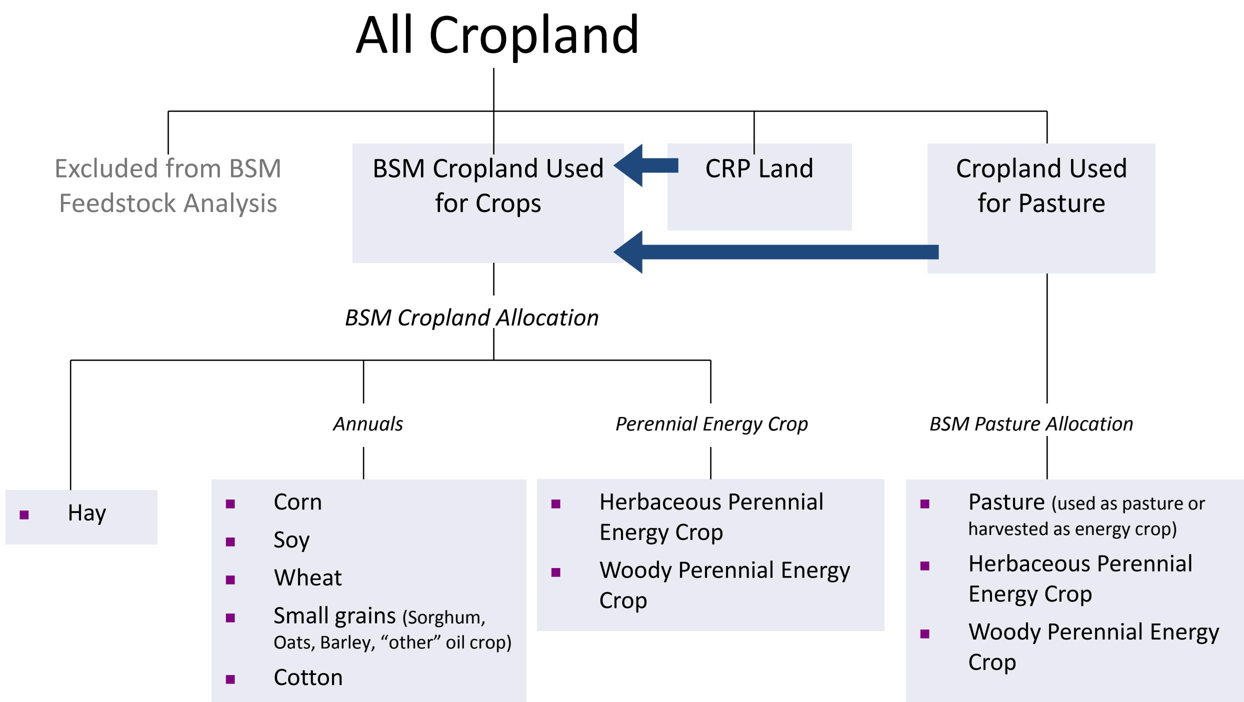

Figure 4. Land categories represented within feedstock supply module

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, but constrained by the requirements of the existing and prospective conversion facilities, as shown in Figure 5.

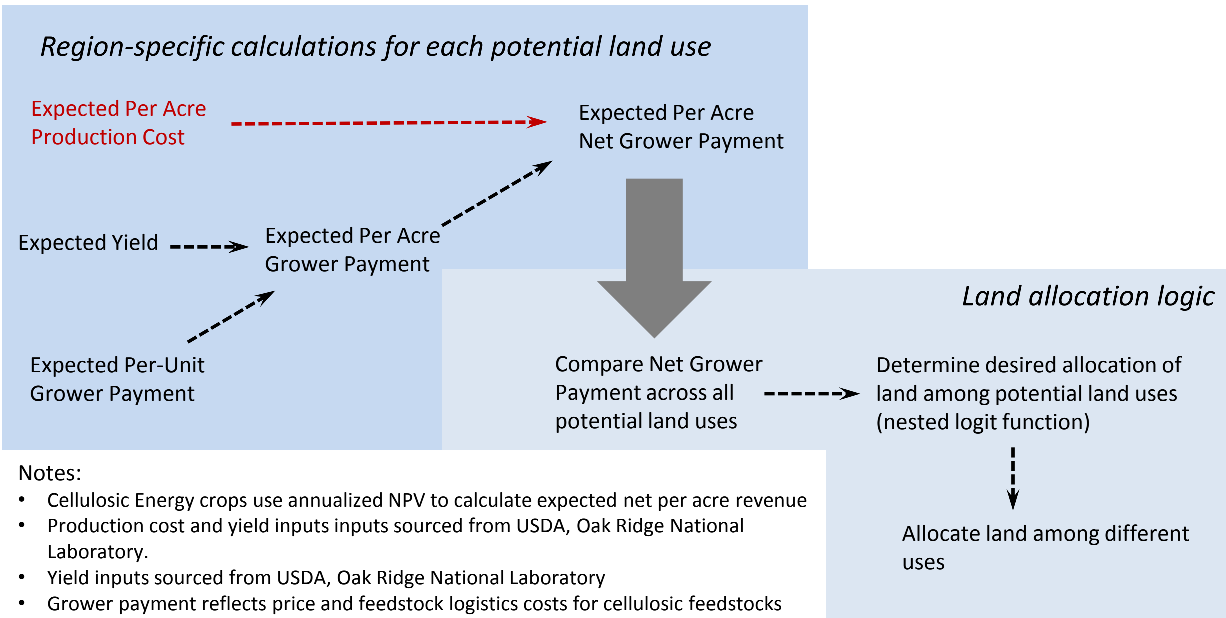

Figure 5. The agricultural land allocation algorithm

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.

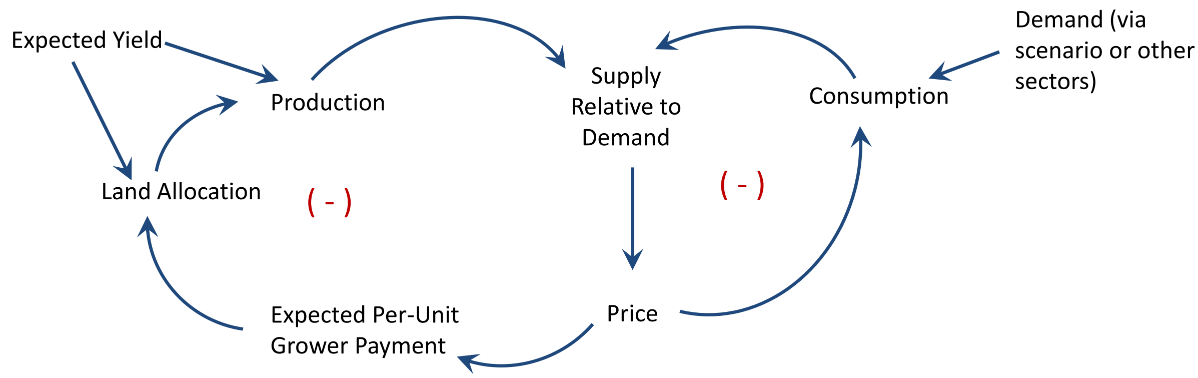

Figure 6. Price feedbacks in feedstock supply module

Feedstock logistics

The feedstock logistics module provides a simple accounting structure that captures the following costs:

- Harvesting and collection
- Transport from "farmgate" to "plantgate"
- Storage, queuing, handling, and pre-processing between farmgate and plantgate.

These costs are used to translate the per-ton price of cellulosic feedstock at the plantgate into a per-ton grower payment at the farmgate. In developing the feedstock logistics module, we have drawn from analyses of the Biomass Logistics Model (BLM) developed at Idaho National Laboratory (INL). In its current form, 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 farmgate to plantgate. 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-to-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.

Conversion sector

The conversion sector is responsible for transforming feedstock into liquid fuels, including ethanol, butanol, and refinery-ready fuels (gasoline, diesel, and 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 scales, 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 net present value.

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
- Incorporation of pre-commercial pilot- and demonstration-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 Aqueous phase reforming	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. As a result we have chosen not to disaggregate these modules by region. The starch-to-ethanol industry, to take another example, is assumed to have reached maturity. Hence, 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 algae 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, as shown in Figure 7, 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.

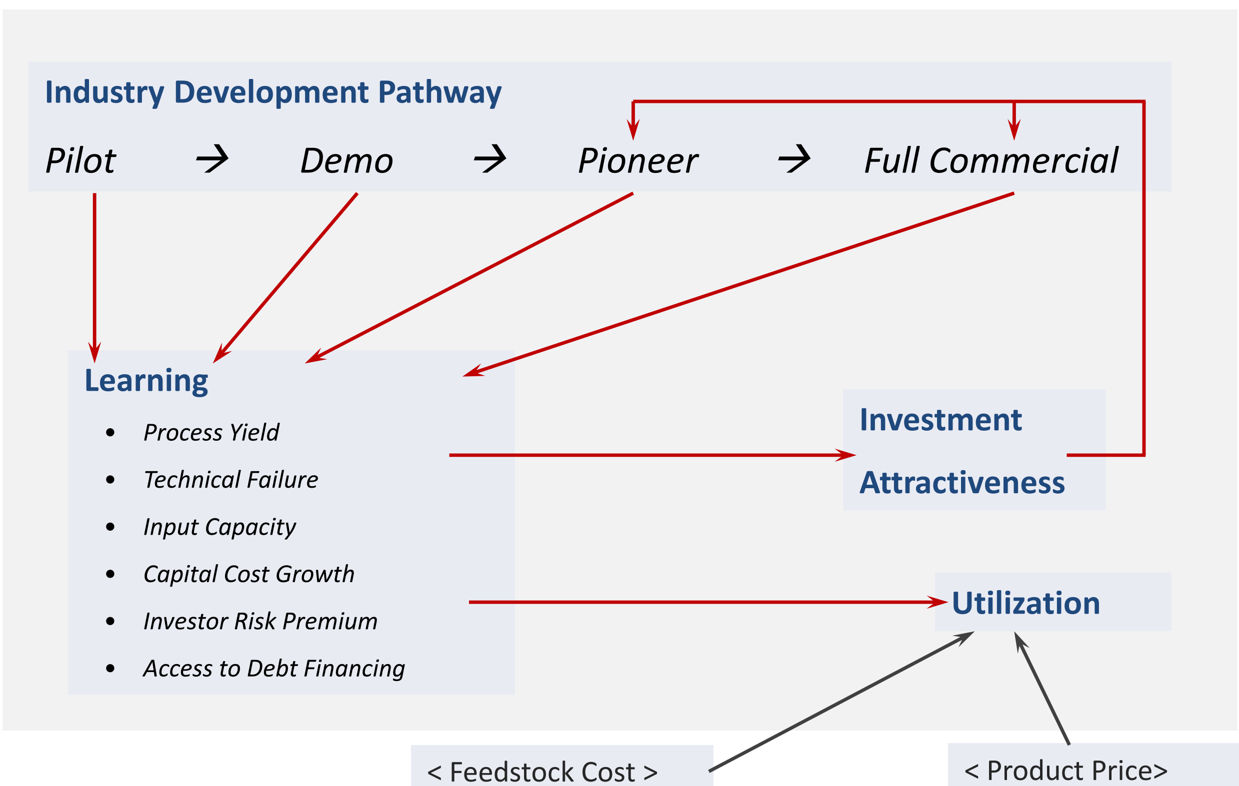

Figure 7. Key interactions within the typical conversion module

Pilot- and demonstration-scale operations

Pilot and demonstration-scale operations are represented with a simple stock/flow structure in the model, as shown in Figure 8.

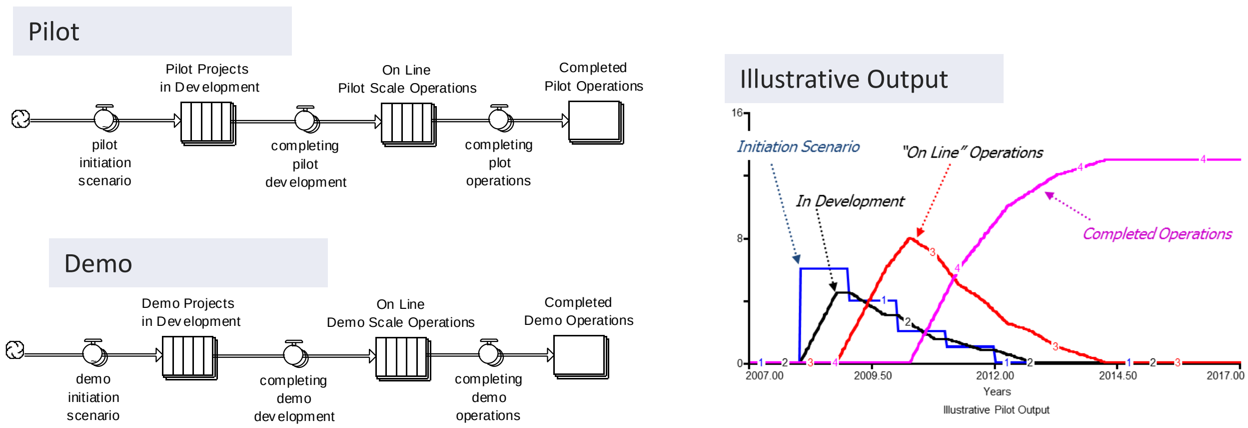

Figure 8. Pre-commercial structure and illustrative output

There are several important features of the pilot- and demonstration-scale structures in the model. First, the structure for each pre-commercial-scale operation is arrayed, based on the number of conversion options at play within the module. Technology on/off switches, set by the user, enable the activation or deactivation of each conversion option. Second, both pilot- and demonstration-scale 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 1/3) 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 (see Table 2 for details). 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.

The co-flow structure here 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 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 "*n*th plant maturity."

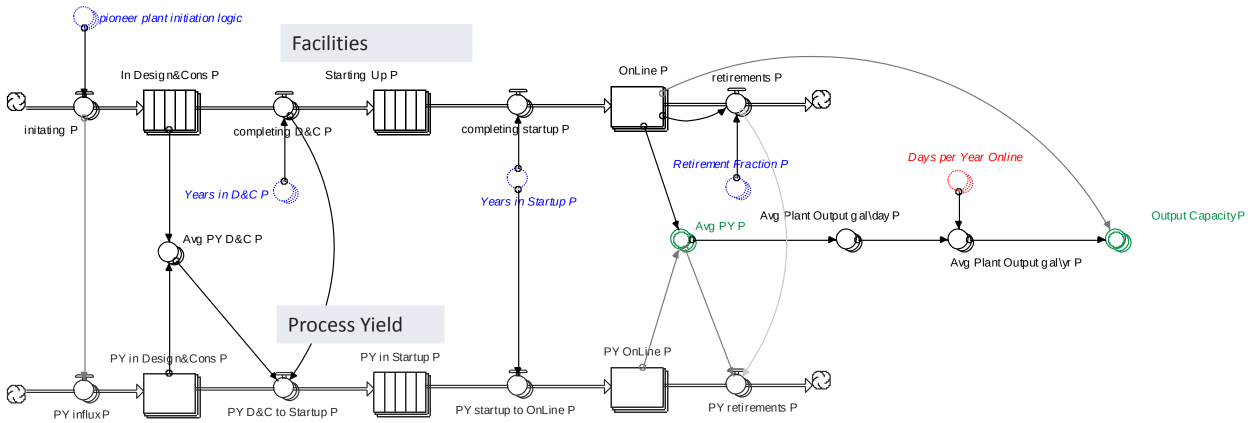

Figure 9. Accounting for pioneer facilities, process yield, and output capacity

Note that structure has been provided to account for retirement of plants. In the current version of the model, we assume a retirement fraction of 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 online 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.

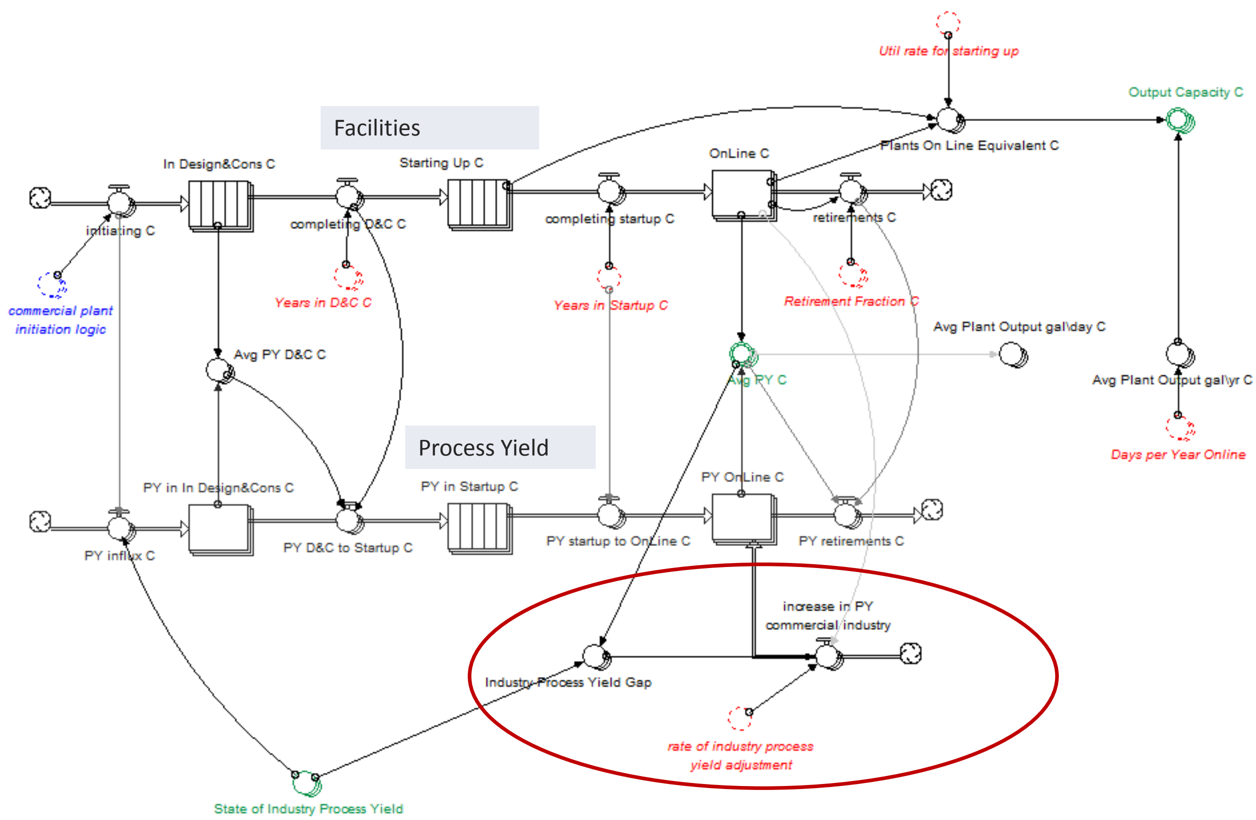

Figure 10. Accounting for commercial operations, process yield, and output capacity

Expected economic value of the "next" investment

In order to represent economic valuation of potential investments, 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. We needed a dynamic economic mechanism to facilitate industry growth, ultimately achieved through 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.

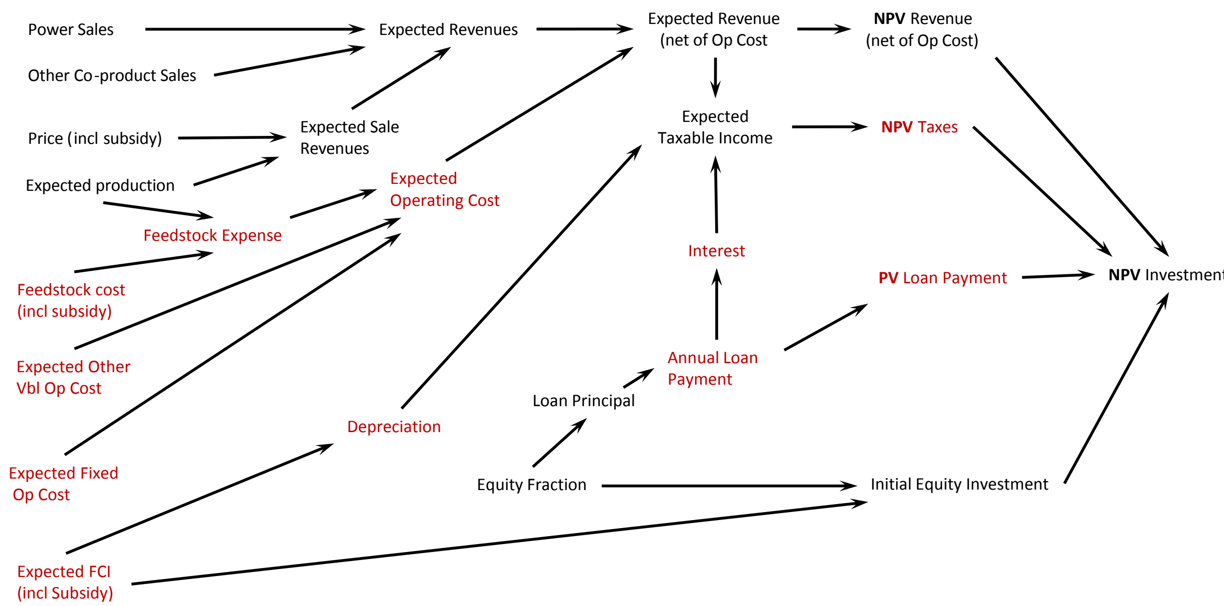

Figure 11. Logic of NPV calculation

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.

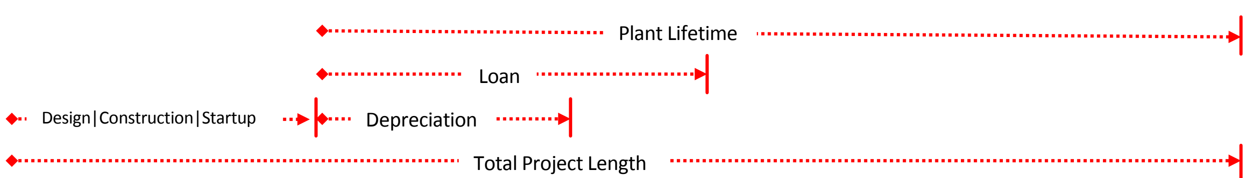

Figure 12. Phases of project life cycle

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, competing across ten regions and often at both pioneer- and commercial-scale. Within the model, each conversion module uses NPV as a basis for determining the attractiveness of the various investment options under consideration using a logit function. (See Appendix B for more details.) The resultant attractiveness metrics are then compared within the relative attractiveness 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 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, as shown in Figure 13.

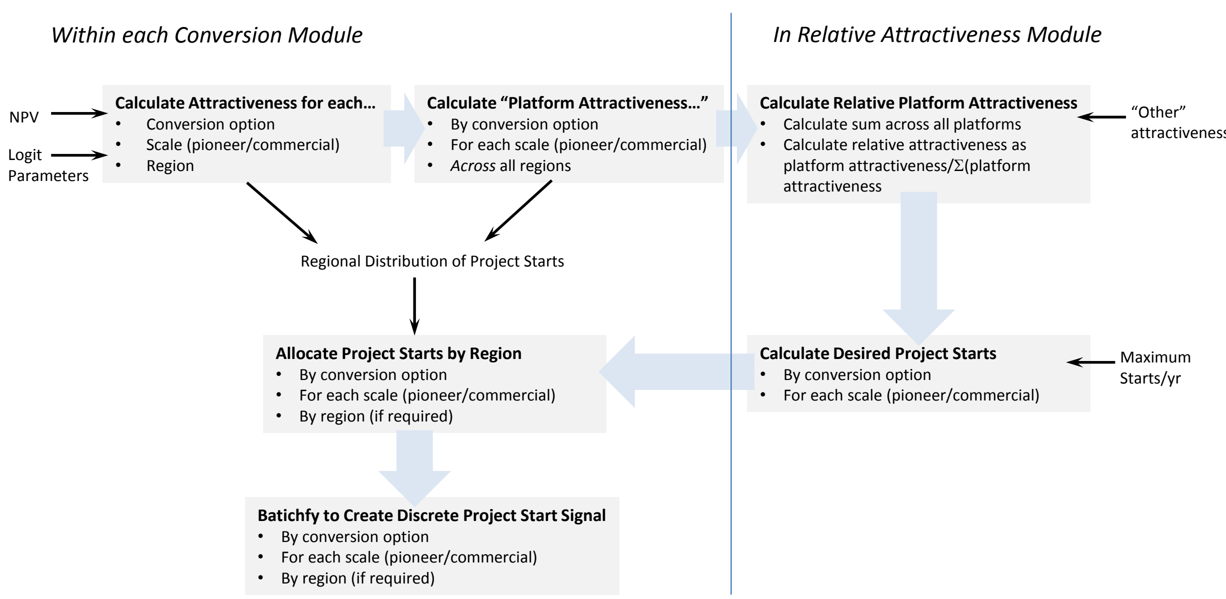

Figure 13. Translating NPV into project start signal

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 "*n*th plant") performance, as reflected by NREL and in 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 & Majluf, 1982) demonstrates the role of accelerating industrial 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, illustrated in Figure 14, 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.
- 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:
    - Process yield
    - Likelihood of "technical failure"
    - Feedstock throughput capacity—the degree to which facilities are able to perform at nameplate capacity
    - Capital cost growth—the premium in capital cost, beyond the *n*thplant estimate, which would be observed if development of a facility was begun today.
    - Investor risk premium—the additional premium, beyond normal hurdle rate, that investors would require for investment in the facility.
    - Access to debt financing—the portion of the expected facility capital cost that would be financed via borrowing (vs. equity investment).

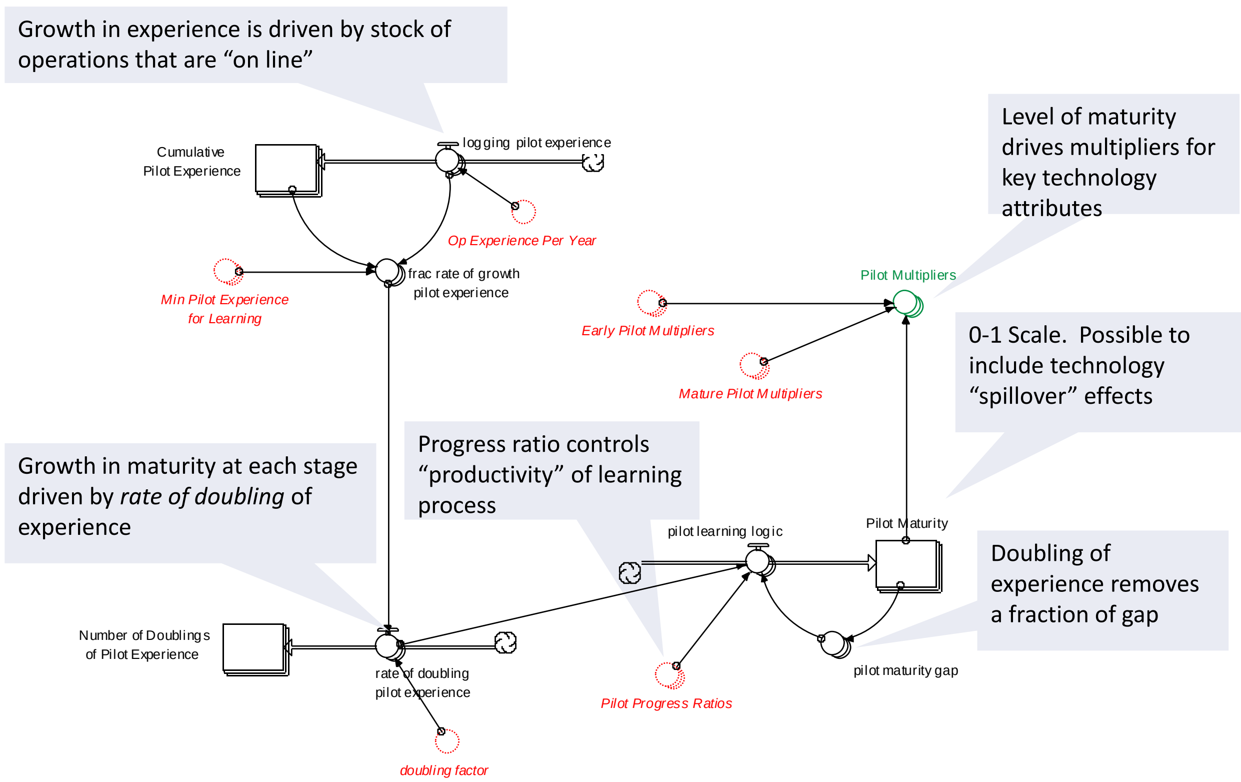

Figure 14. Industry learning curve structure

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.

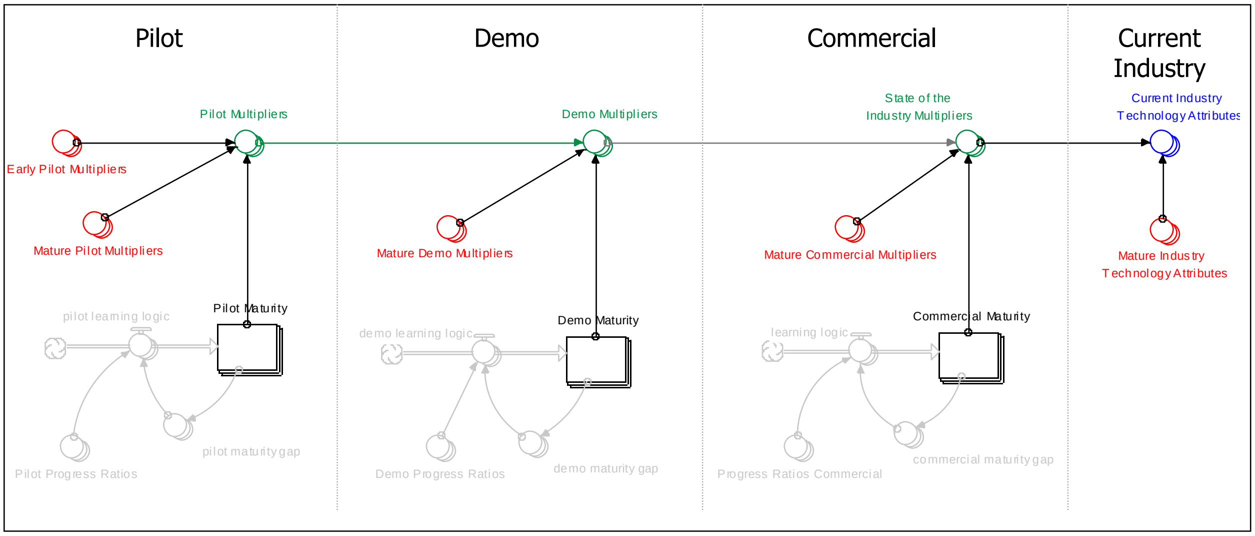

Figure 15. Cascading learning curves

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. Figure 16 illustrates that 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.

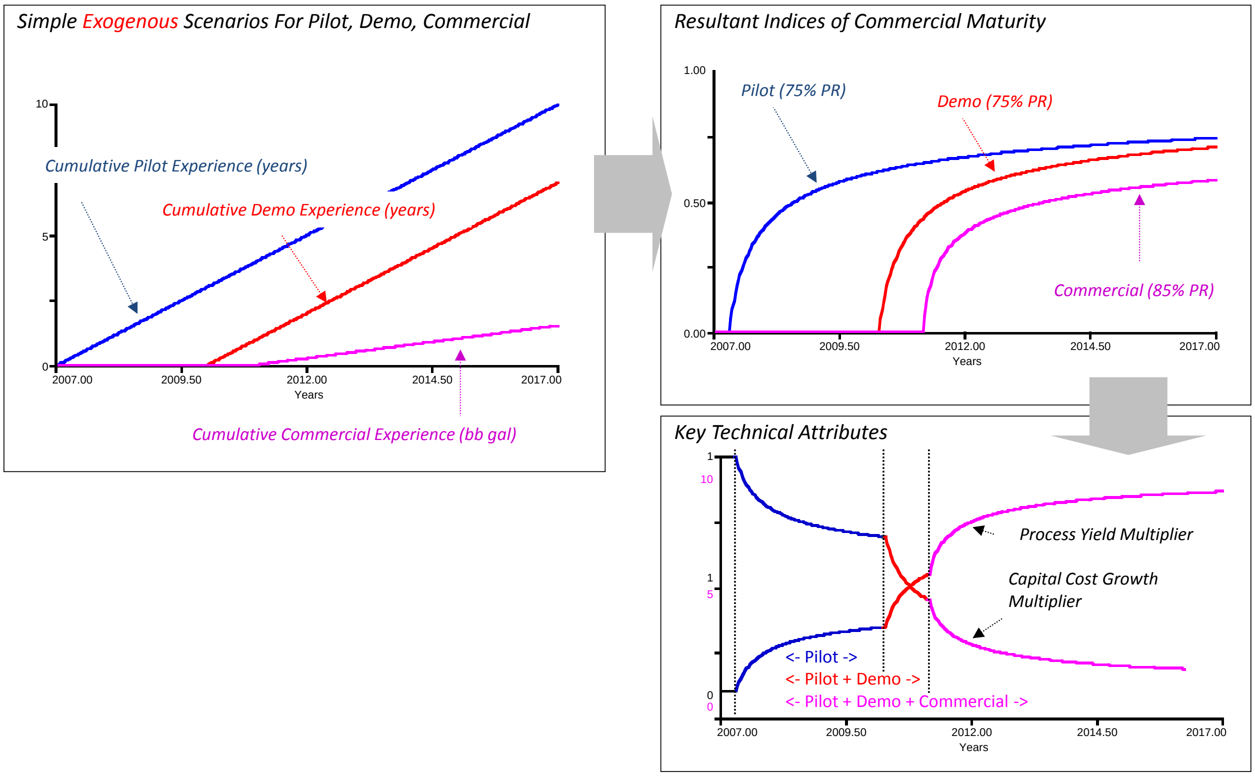

Figure 16. Illustrative learning curve dynamics

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" (as shown in Figure 17). Each of these mechanisms is a positive feedback loop, reinforcing development of the conversion option in question.

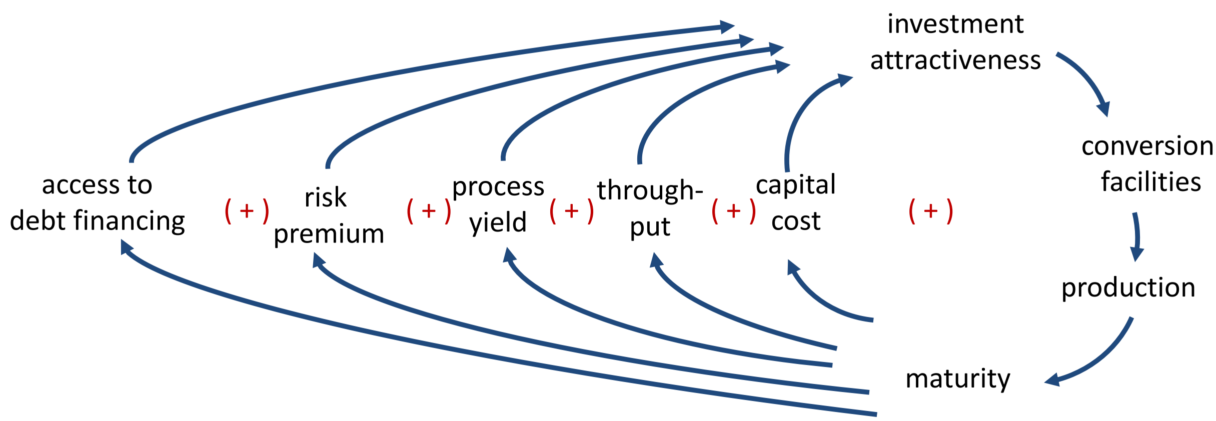

Figure 17. Key feedbacks emerging from learning curve structure

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, as shown in Figure 18.

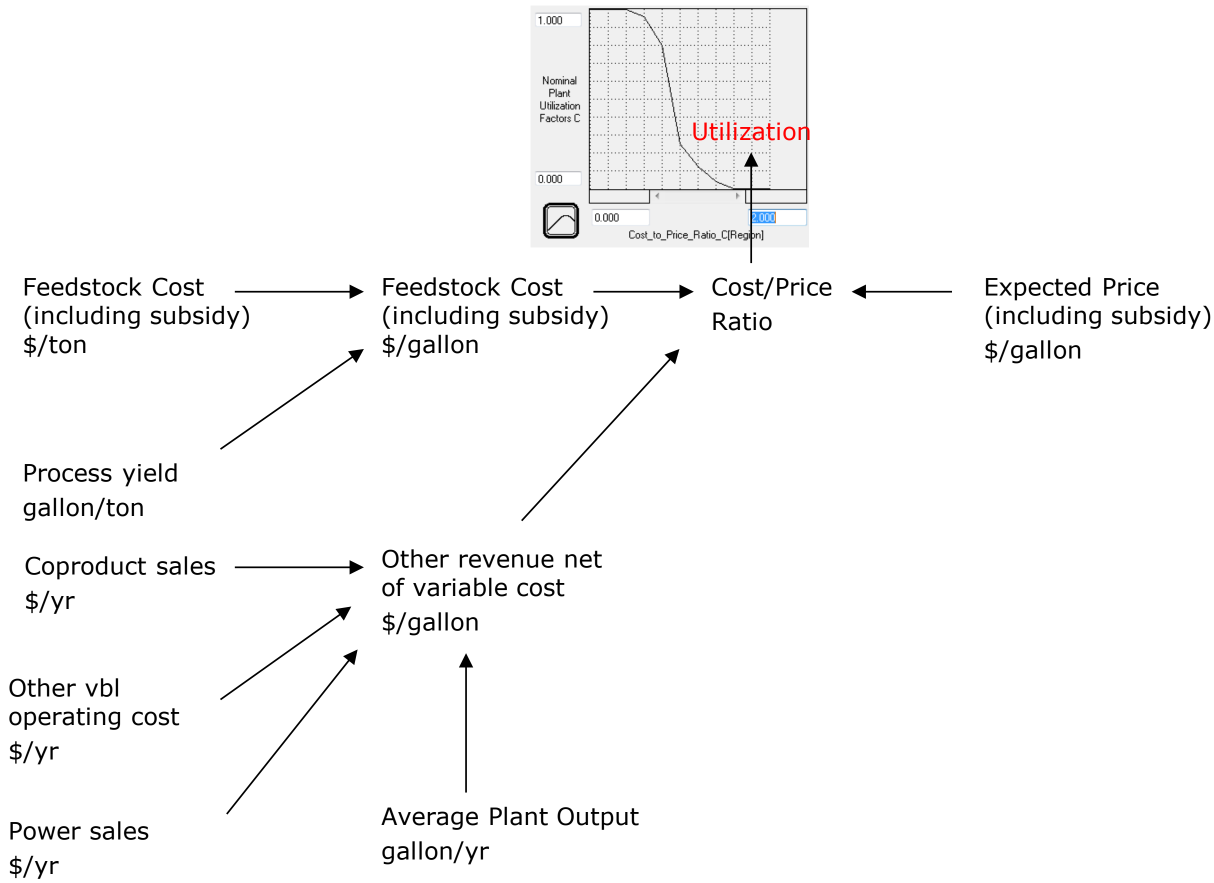

Figure 18. Determining utilization from cost-price ratios

Utilization is a central element of the feedback structure within BSM, as it controls both the production of products and consumption of feedstock.

Conversion sector 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, and 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 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.

Downstream ethanol sector

The downstream ethanol 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 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.

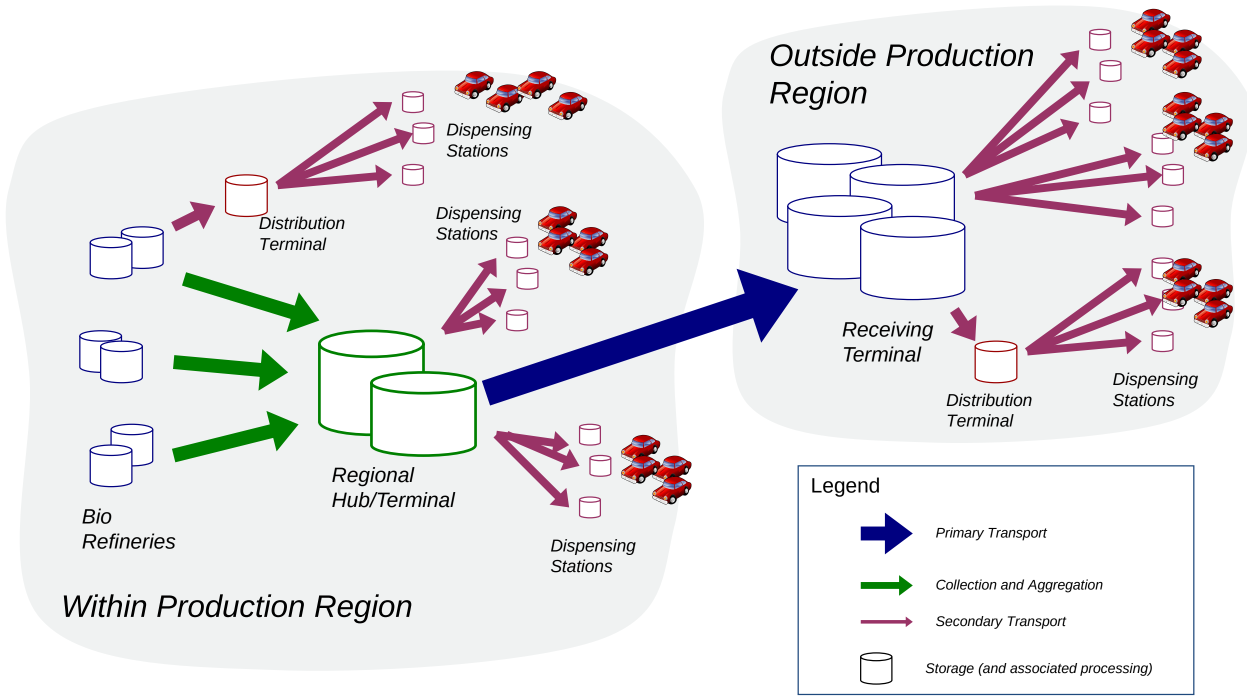

Figure 19. An overview of downstream dynamics

To support these dynamics, multiple modules comprise the downstream sector of BSM[3], including:

- Distribution logistics
- Dispensing station
- Fuel use
- Pricing and inventory.

Distribution logistics module

A fundamental challenge 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

[3] Detailed analysis of downstream ethanol dynamics can be found in (Vimmerstedt, Bush, & Peterson, 2012).

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, time delays, and feedback loops associated with regional build-out, as shown in Figure 20.

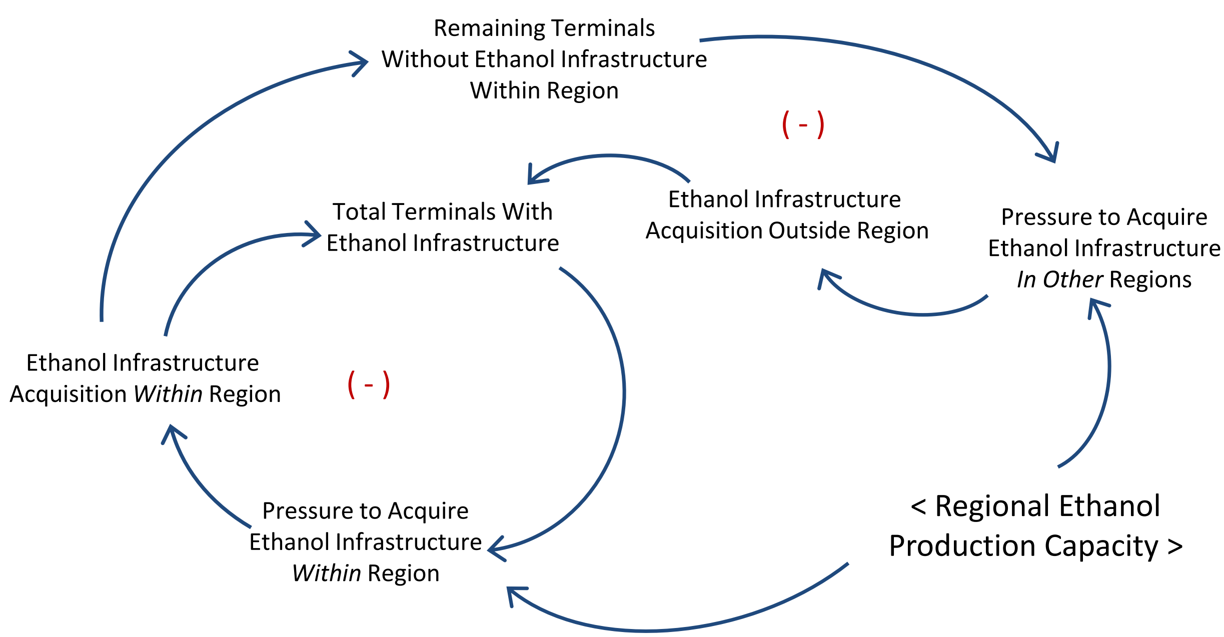

Figure 20. Distribution logistics

A two-stage supply-push approach (first within a region, and then across regions) is embedded within the module. Within a region, the model first 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 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 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 hypermarts. The fundamental decision for each station is the acquisition of tankage and dispensing equipment required to dispense high-ethanol blends into FFVs. The module assumes that 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 high-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 NPV 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 (Figure 21). The second provides a detailed view into the NPV calculation that undergirds the decision to invest in high-blend tankage and equipment (Figure 22).

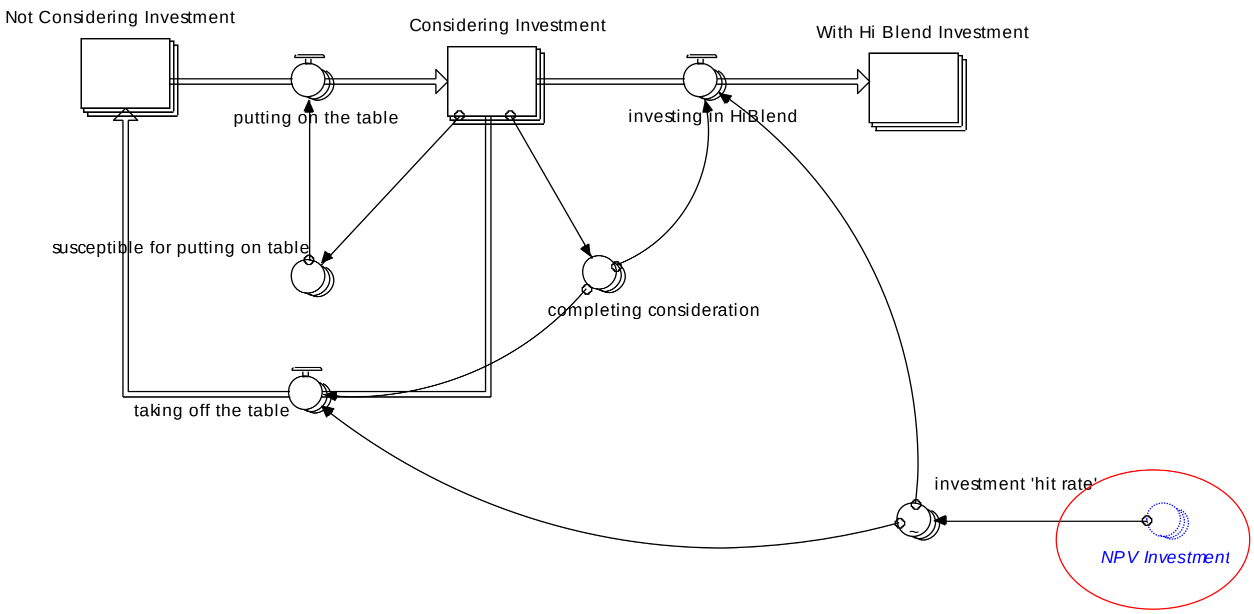

Figure 21. Dispensing station accounting structure; NPV calculation captures estimated costs and revenues of prospective investment

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, a portion of those stations not considering investment put the investment decision on the table each year. Based on the economic viability of the investment (as reflected in the NPV of the 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 repurpose versus new investment (to account for different capital costs associated with repurposing versus new investment in tankage and equipment).

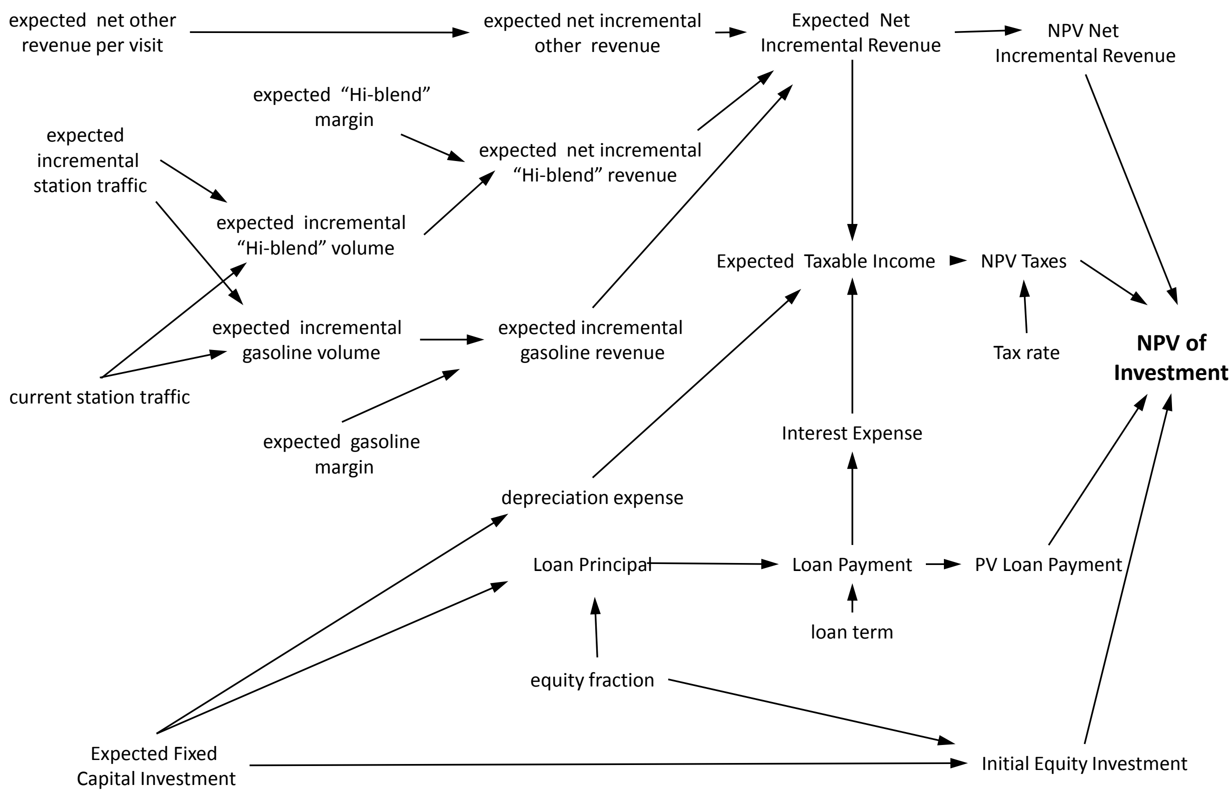

Figure 22. Logic behind NPV calculation for stations

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 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 FFVs 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 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, as shown in Figure 23. The first component accounts for the affinity of FFV owners toward high-blend fuels. The second uses a logit function (see Appendix B) to allocate fuel use between for FFV owners who are "occasional" and "regular" users of high-blend fuels.

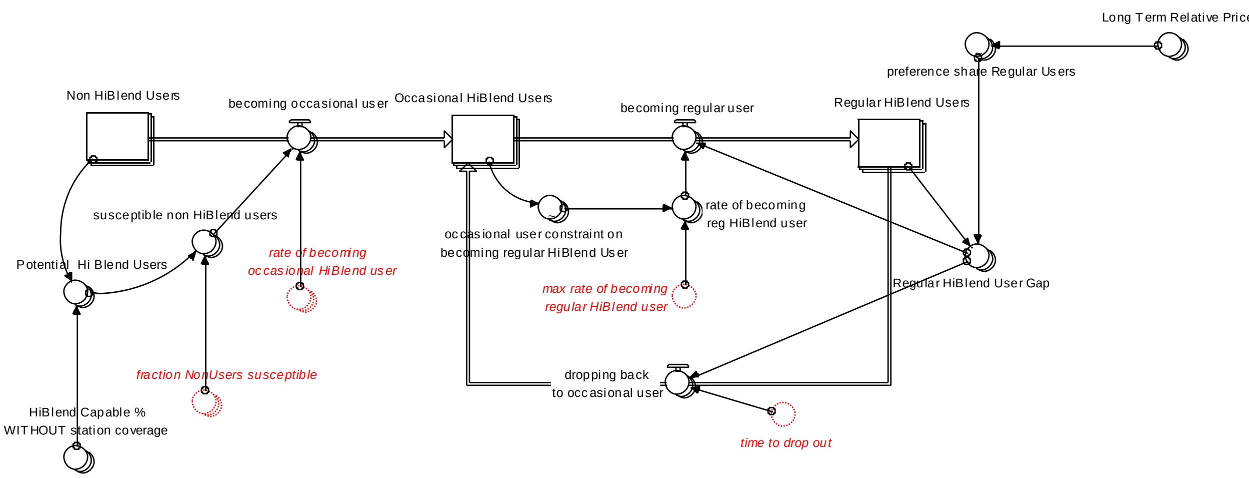

Figure 23. FFV accounting structure

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 multiplied against potential high blend fuel consumption in order to generate actual high blend consumption within each region, as shown in Figure 24.

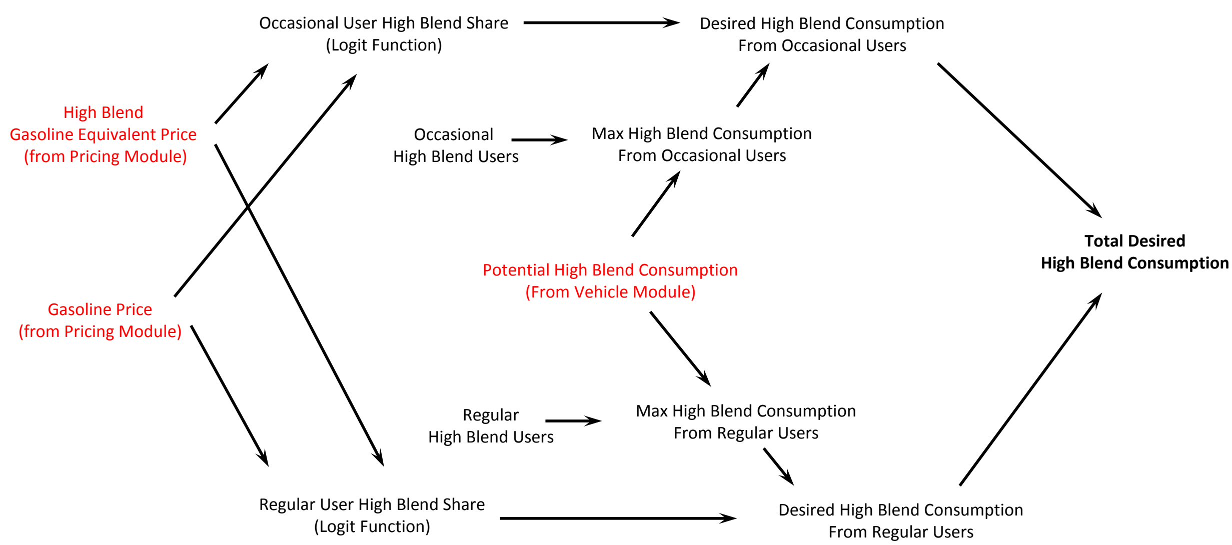

Figure 24. Logic behind high-blend consumption

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 here.

Ethanol inventory is aggregated across the entire supply chain within each region, allowing for cross-regional movement of ethanol based upon regional surpluses or shortfalls within each region (as shown in Figure 25).

There are several important features to this pricing and inventory structure. 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, reflecting total 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.

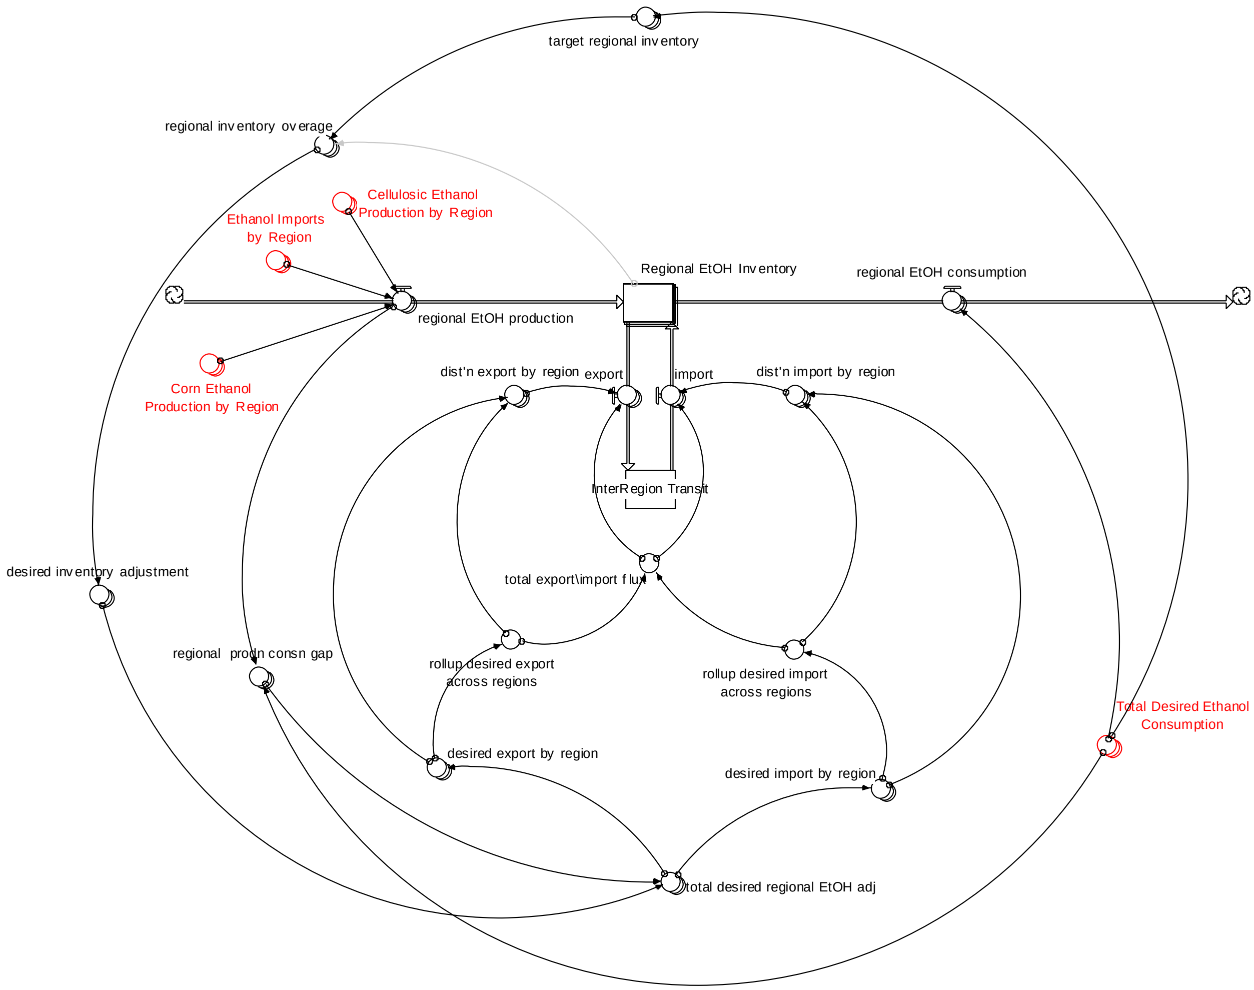

Figure 25. Downstream ethanol inventory dynamics

This cross-regional movement algorithm is relatively straightforward, and as described 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 (see 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 26, 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.

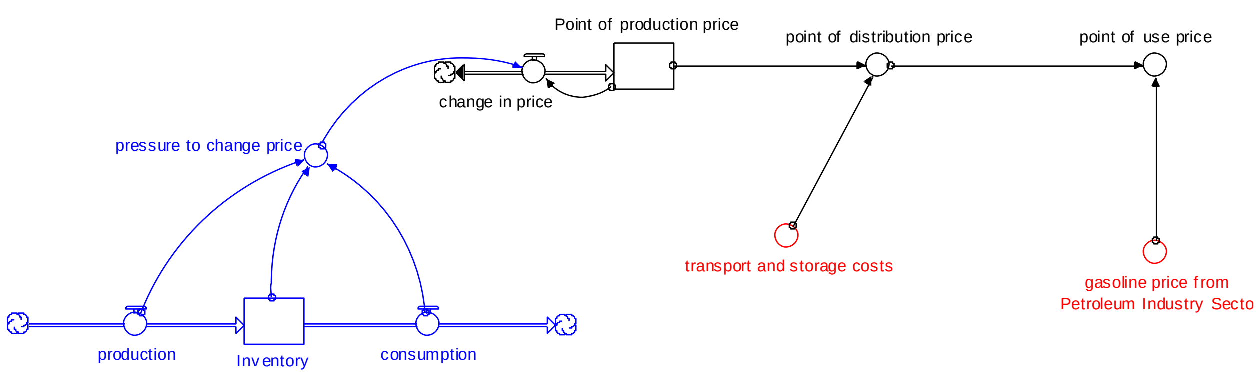

Figure 26. Simplified ethanol pricing structure

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 oxygenate 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/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.

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. 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, and vehicle mortality. The model aggregates vehicles nationally. Regional population distributions 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 (automobiles and light trucks) and 10 engine types (gasoline, diesel, plug-in hybrid (PHEV), hydrogen, compressed natural gas, FFV, gas hybrid electric, gas PHEV, bi-fuel, and other) within the light duty fleet. For each of these 20 combinations, a scenario for new-vehicle sales over time is accompanied by a scenario for new-vehicle efficiency. The model dynamics track the implications of these new vehicle scenarios for overall vehicle efficiency and resultant fuel demand, as shown 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.

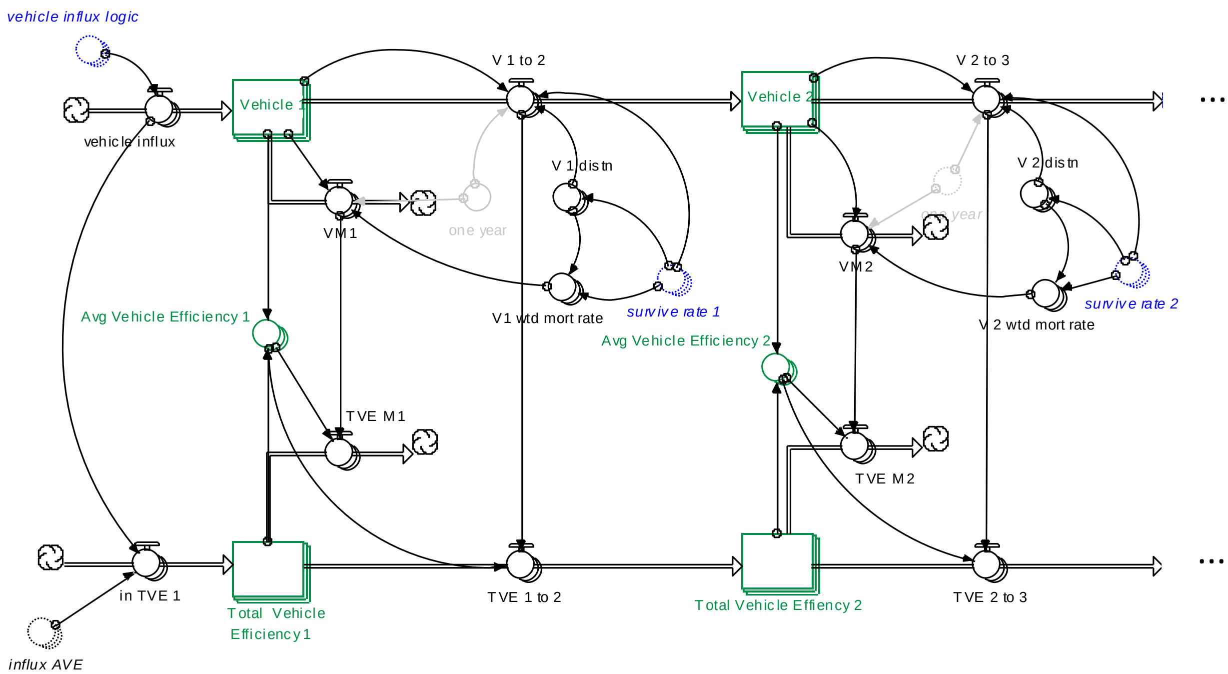

Figure 27. Tracking vehicles and efficiency (2 cohorts)

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 other similar items. 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 vehicleinflux side, levers within the vehicle module enable use of (or departure from) Annual Energy Outlook (AEO) projections for inflow volume, inflow mix, and efficiency. Similarly it is possible to use or depart from AEO mortality rates, and to use or modify AEO scenarios.

Oil industry sector

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

- A set of scenarios used to determine crude oil prices

- Refinery product prices for diesel, jet fuel, and gasoline
- Algebraic relationships that translate crude oil prices, refinery product prices, and an assumed refinery "drop-in point" for each infrastructure-compatible pathway into price inputs for the different conversion modules and for the downstream pricing and inventory module
- Accounting structure that captures petroleum displaced by diesel, jet fuel, and gasoline produced by the different infrastructure-compatible pathways.

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. 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, an 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.

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 (plantgate)			
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 consumption from FFV • Potential lo blend consumption from non-FFV • Potential hi-blend consumption • Potential gasoline consumption

Data inputs

Multiple data inputs are required to run BSM, including agricultural cost and yield parameters for the feedstock module, performance and learning parameters for the various conversion modules, logit coefficients, petroleum prices, and 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, shown in Table 4.

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 Pacey report
	Yields	ORNL/POLYSYS	yield growth treated as assumption/scenario
	Calibration data for production, prices	USDA baseline (United States Department of Agriculture)	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 (RW 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), 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 Ethanol	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 & 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	
	Logit parameters for fuel	Assumption	

	use		
	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

Analysis infrastructure

As has been outlined above, the BSM is a robust model that has undergone rigorous testing, validation and refining by the BSM team, and was built to explore multiple facets of the biofuels supply chain and its numerous drivers, bottlenecks, and system interactions. The model was designed to be a comprehensive, agile tool that would allow U.S. Department of Energy (DOE) to perform quick-turnaround analyses in response to evolving policy, scenario, and research questions. In order to quickly perform multiple runs of the model with different scenario inputs and to always be able to review runs that were made historically, it was important to set up a framework for storing all inputs and outputs of the model for all runs made for important analyses.

We use modern software-engineering methodologies to maintain model quality and enable flexibility and responsiveness in response to analysis requirements that evolve as new bioenergy issues gain interest from stakeholders. 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 model. Input data are stored, raw data sources are archived, and provenance/pedigree metadata is tracked within a relational database: furthermore, input data sets are processed within that 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 database-centric approach makes it easy to develop and package "scenario libraries" for stakeholder use, as shown in Figure 28.

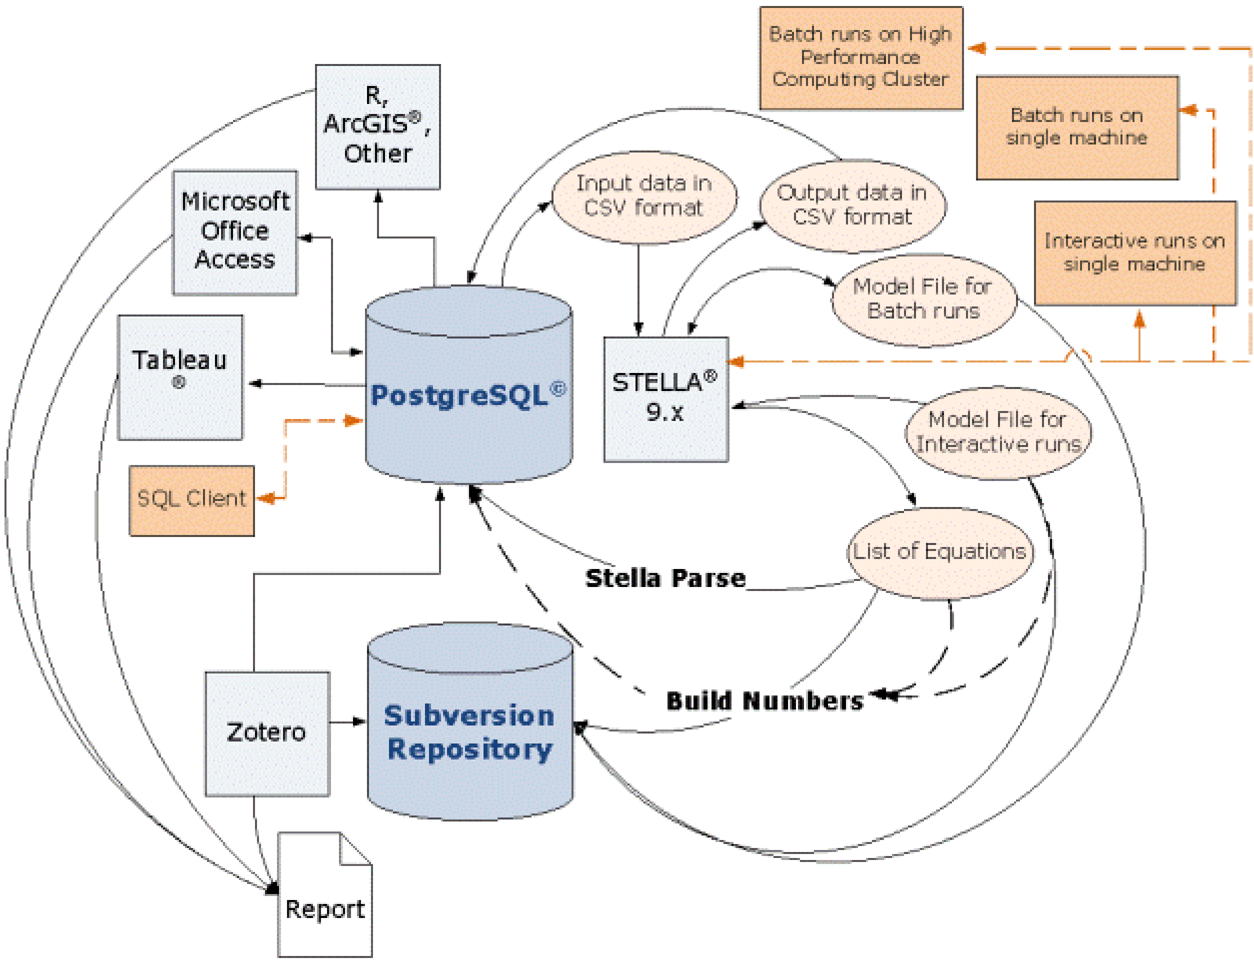

Figure 28. Computing infrastructure for BSM.

The aforementioned computing infrastructure supports a high-throughput analysis process that is outlined in Figure 29. In particular, it enables 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 output into the database. To further enhance the approach, the BSM source files can be copied and run on multiple machines at once to quickly make thousands of runs simultaneously. The required output variables are specified in the database, and the values for these variables are taken from STELLA for the specified simulations and transferred to the database. In this way, we have a central system where any team member can re-create any past scenario—either by viewing the previous runs or finding the correct model on the model repository. The outputs can then be imported into any graphics software to visualize the simulation results, analyze trends, and develop insights. Using these techniques, we have been able to analyze and compare thousands of BSM runs with little effort, completing analyses that have been included in over 15 internal analysis reports, 8 external publications, and 10 forthcoming publications.

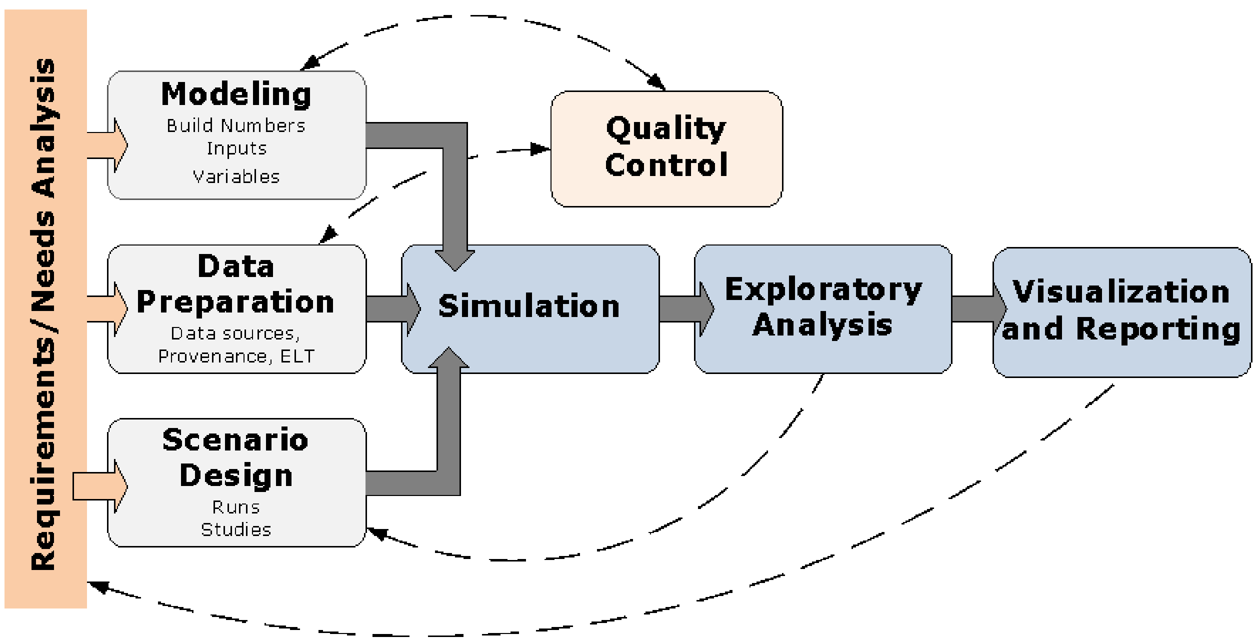

Figure 29. High level overview of the BSM analysis process.

Analysis, scenario development, and insights[4]

Specific policy-relevant scenarios or past scenarios can be used to drive BSM simulations, though BSM is not limited to scenario analysis. Under a specified scenario, BSM can be used to track the hypothetical development of the biofuels industry given the deployment of new technologies within various elements of the supply chain and the reaction of the investment community to those technologies and given the competing oil market, vehicle demand for biofuels, and various government policies over an extended timeframe. Note, however, that high-level models such as the BSM are not typically used to generate precise estimates but rather to:

- analyze and evaluate alternate policies
- generate highly cost-effective scenarios
- identify high-impact levers and bottlenecks
- focus discussion among policymakers, analysts, and stakeholders.

When BSM output includes unexpected system behaviors, modeling assumptions—particularly the behavioral aspects of decision making and the adequacy of the representation of feedback—need careful reexamination to distinguish potential insights from model limitations. The model itself often indicates what assumptions need the most scrutiny; hence, it helps define the research and learning agenda.

[4] Analysis efforts using BSM have been ongoing since early in 2010, initially using an earlier version of the model that focused on ethanol from cellulose and starch crops. Beginning in the fall of 2011, the BSM team has been designing experiments, creating scenarios, and conducting analyses using the current version of the model which includes both ethanol and infrastructure-compatible fuels.

Although BSM inputs can be altered to include any combination of policies, initial analysis efforts included establishing a "reference policy case" to which subsequent scenarios could be compared. The BSM reference policy case includes moderate incentives for ethanol production and a 50 cent per gallon gasoline tax (which could be interpreted as a "carbon" or emissions tax in dollars per ton of carbon dioxide). Policies are phased out in a staged manner, with the policies involving grants for capital equipment or loan guarantees ending earlier and the policies involving volumetric subsidies phasing out anywhere from 2020 to 2050. Each of the policies included in the reference case is based on historical precedence or future plausibility.

As BSM functionality increased and research questions from stakeholders became more sophisticated, we created an expanded list of scenarios to be easily incorporated, tested, and analyzed. These scenarios are not intended to be prescriptive or comprehensive, but instead represent an extended backdrop of cases against which policies can be tested and possible industry evolution can be explored. The scenario library approach has proven to be quite agile and useful from an analysis perspective, and the BSM team expects this functionality to be augmented and expanded as part of the project's ongoing development efforts.

Table 5. Scenarios in current BSM scenario library

Scenario Name	Scenario objective/constraints	Strategy employed
1: Minimal Policy	Starch until 2012	Apply minimal subsidies and policies
2: Ethanol Only	Ethanol pathways only	Provide support for ethanol only; analogous to BSM reference case
3: Equal Access	All pathways in order to produce 36 billion gallons/year by 2031	Allow all fuel types equal access to generous scenario subsidies
4: Output-focused	To maximize growth restricted to $10 billion per year	Target most promising technology and withhold subsidy access from other pathways
5: Pathway Diversity	To maximize pathways restricted to $10 billion per year	Design subsidy timeline to enable take-off of multiple fuel pathways by staggering start and end dates based on pathway progress and potential

Scenario 1: Minimum policy

The minimum policy scenario (Table 5, Figure 30) includes only a $0.45 price subsidy at the point of production for starch ethanol that expires in 2012; it does not have any additional subsidies directed towards renewable fuel production. Without government intervention in the form of renewable fuel subsidies and given the oil price assumptions used in the model, neither the cellulosic ethanol nor infrastructure-compatible fuel industries gain industrial momentum, and thus fail to "take off" to any significant extent. Starch ethanol is able to satisfy the market for oxygenate in gasoline. The declining demand for oxygenate over the BSM time period is attributable to the overall decline in gasoline demand as more fuel-efficient vehicles enter the market in response increased Corporate Average Fuel Economy (CAFE) standards.

Scenario 2: Ethanol only

The ethanol only policy, (Table 5, Figure 30) applies all subsidies to the renewable ethanol industry exclusively. With all cellulosic feedstock available to the ethanol industry, the cellulosic ethanol industry is able to reach nearly 9 billion gallons in annual production (bgy). Annual spending peaks at $6 billion (aside from the initial starch subsidy.)

Scenario 3: Equal access

The renewable fuel standard (RFS2) mandates that 36 billion gallons of renewable liquid transportation fuels will be in the market place by the year 2022; the annual RFS2 volumes are allowed to be adjusted by the U.S. Environmental Protection Agency (EPA) based installed capacity and the amount of fuel demanded (US EPA 2011). The equal access scenario is designed to mimic pathway agnostic policies such as the RFS. In this scenario, subsidies are set to levels that spur renewable fuel output to RFS2 levels (i.e., 36 bgy). The results of this scenario need to be viewed in the context of the initial settings and assumptions regarding the industrial maturity of the infrastructure-compatible fuel technologies examined.

Scenario 4: Output-focused

Scenario 4 focuses subsidies on one pathway, i.e., fast pyrolysis. Fixed capital investment subsidies and loan guarantees were limited to fast pyrolysis and not available to other fuel pathways (Figure 30). Spending was limited to $10 billion per year, and after the expiration of the starch ethanol price subsidy at the end of 2012, total subsidies reached a peak of only $5.3 billion in 2023. As the output for fast pyrolysis grows in later years, the fast pyrolysis subsidies grow as well because of the price subsidy on each gallon of fuel. The exposure to loan guarantees is not counted in the total subsidy figure. Infrastructure-compatible fuels – almost exclusively fast pyrolysis – contributed 34.1 billion gallons to the total 51.4 billion gallons produced in 2030.

Scenario 5: Pathway diversity

In the "pathway diversity" scenario, we explored the possibility of promoting pathway production diversity by launching four different technologies to produce volumes of significant output (over 1 billion gallons) with a total annual budget of $10 billion. After selecting the four most competitive technologies through preliminary analysis (Fischer-Tropsch, fast pyrolysis, fermentation, and methanol-to-gasoline), different subsidy amounts, start times, and durations for each technology were applied in order to achieve output levels spread most evenly across the technologies. The staggered start times and durations increase the attractiveness of technologies with promising mature commercial-plant techno-economic parameters but starting with limited industry maturity and experience. By the time pathways are ready to take off (around 2023), technologies are on more even footing, allowing for greater pathway diversity than the other scenarios. Industry output for infrastructure-compatible fuels reaches the 3 -billion-gallon volumetric limit just before 2023, at which time heavier "startup" subsidy values switch to lower "background" subsidy values (as indicated in Figure 30). Annual production reaches 34.9 billion gallons of renewable fuels per year in 2030 and reaches a peak of $8.9 billion of spending in year 2023. Of the total production in 2030, fast pyrolysis, Fischer-Tropsch, fermentation, and methanol-to-gasoline produce 5.7 billion, 5.3 billion, 1.3 billion, and 5.5 billion gallons, respectively.

Scenario insights

BSM simulations based on these scenario libraries have provided a wide range of insights to the project team and policy-makers alike. Potential policies designed to accelerate the development and sustainability of the biofuels industry can be easily tested across these embedded scenarios under a wide range of assumptions regarding the magnitude, duration, and sequencing of various policy interventions. Subsequent sensitivity studies on important model parameters are used to quantify the responsiveness of various key BSM output metrics to policy initiatives, alone or in combination, in these different scenario cases. This rigorous testing has built confidence in the robustness of BSM, as well as informed key insights into the nature of the evolution of the biomass-to-biofuels supply chain.

Insight 1: Momentum in the infrastructure-compatible fuels industry causes hierarchical competition for feedstocks, thus reducing ethanol market share.

The existence of three renewable-fuel industries creates an interesting hierarchical competition: the infrastructure-compatible fuels industry competes with the cellulosic ethanol industry for feedstocks, while the starch ethanol industry competes with the cellulosic ethanol industry for market share (both low- and high-blends). In the minimal policy scenario, the unsupported biofuels industry produces only 12 billion gallons of starch ethanol output in the year 2030; the cellulosic ethanol industry produces only about 60 million gallons of ethanol in the year 2030. Starch ethanol is well established and continues to provide oxygenate for gasoline and high-blend ethanol gasoline (E85), accounting for most, if not all, of the market for ethanol. Without government intervention, the starch industry does not face competition from the cellulosic ethanol industry and is able to meet all the ethanol demand.

When subsidies are applied exclusively to ethanol (with an emphasis on cellulosic ethanol), the hierarchy between starch and cellulosic ethanol is salient. To prevent significant industrial bottlenecks and encourage market penetration, downstream infrastructure subsidies are critical. When cellulosic ethanol subsidies are high, the industry takes market share away from the starch ethanol industry, but the latter is able to recover in the long-term because of its maturity. Although this case subsidizes cellulosic ethanol heavily relative to all other pathways, ultimately growth is restricted by limited market for ethanol (described above) as starch ethanol and cellulosic ethanol compete for the same market. In the competition for market share, the minimal policy and ethanol only scenarios confirm that cellulosic ethanol is able to compete with starch only with sufficient subsidies in the developing years.

Providing subsidies for all pathways in the RFS2 scenario, the infrastructure-compatible fuels industry is able to outbid cellulosic ethanol, driving up feedstock costs, which disadvantages the cellulosic ethanol industry compared to the starch ethanol industry. Additionally, unlike ethanol, the infrastructure-compatible fuels are modeled with potentially unlimited demand and no interference (bottlenecks) from lack of downstream infrastructure. The net effect of these factors is a 59% reduction in cellulosic output by the end of the simulation (i.e., 2030) relative to its peak output, and a rapidly growing infrastructure-compatible fuels industry.

Insight 2: RFS2 volumes are achievable in 2030 with heavy startup subsidies.

Under the ethanol only scenario (Table 5), a total of 35.9 bgy of renewable fuels is produced in the year 2030. Infrastructure-compatible fuels contribute over half of this amount (18.7 billion gallons), while the starch- and cellulosic-based ethanol industries cumulative comprise 17.2 bgy. The RFS2 timeline is shown to be impractical because the high-blend ethanol market applies pressure on the system (i.e., the high-blend ethanol market is not large enough), and the infrastructure-compatible fuels industry is not mature enough to produce 36 billion gallons in 2022. Reaching this level of production requires investment in the form of start-up subsidies, particularly fixed capital investment subsidies and loan guarantees for commercial-scale facilities. Total annual subsidies for the industry peak at a demanding $34.2 billion in the year 2024. The start-up fixed capital investment and loan guarantee subsidies (especially commercial-scale) are more effective at quickly building the industry than other subsidies but are far more costly than the others.

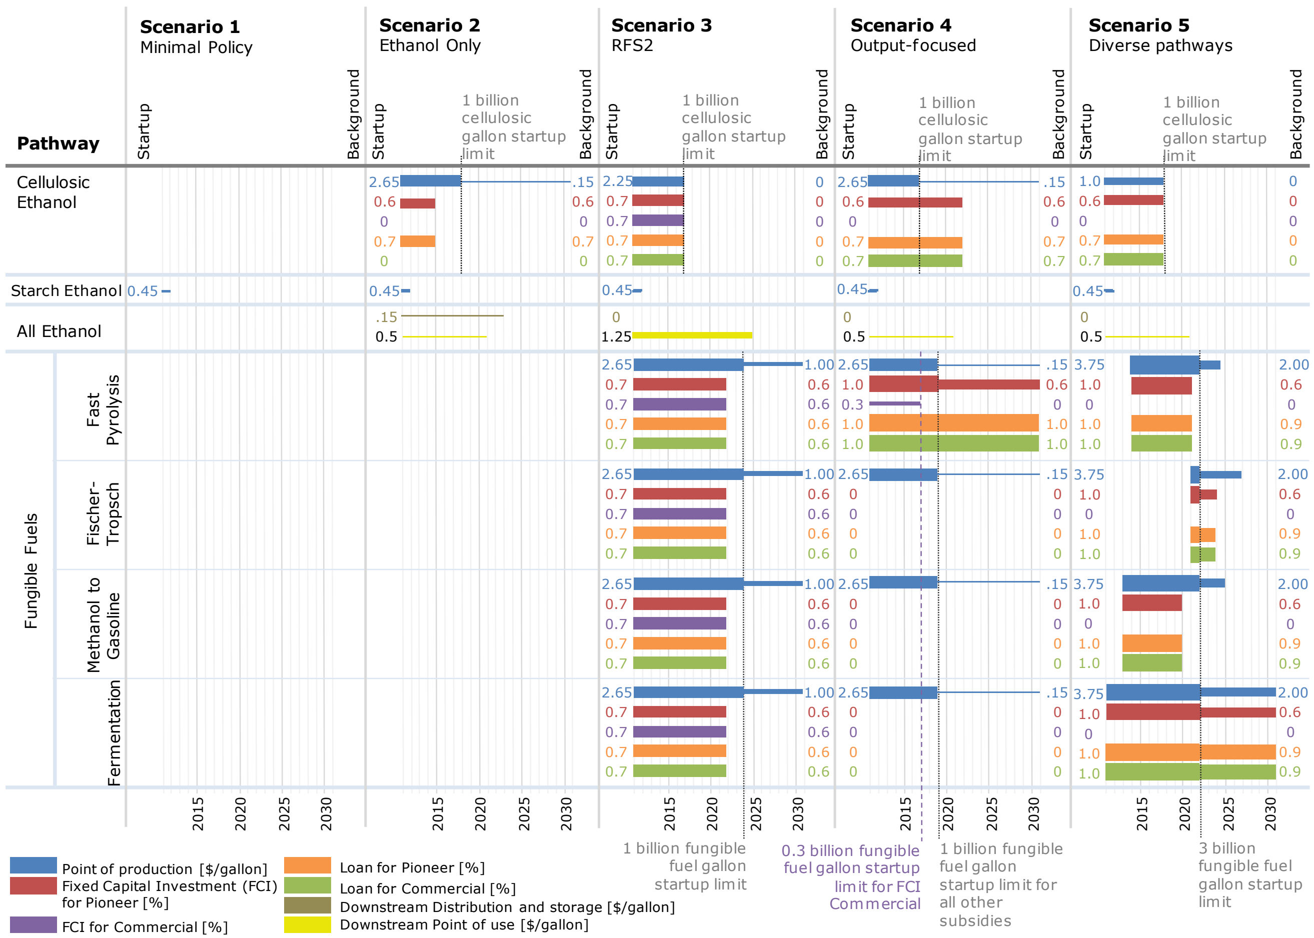

Figure 30. Subsidy summary for scenarios 1 through 5. "Startup" value refers to values left of the annotated limit line(s) for the technology, background values refer to values right of each respective line; thickness of the duration bar indicates relative magnitudes of subsidies. Start and end years can be mapped to the year on the x-axis.

Insight 3: Production levels can exceed RFS2 levels if subsidies promote the most economically attractive pathway

Even though the volumetric output of Scenario 4 was higher than Scenario 3 (RFS2), the spending in Scenario 4 was less than that of Scenario 3 (RFS2). In Scenario 3, most years had annual spending on subsidies exceeding $10 billion, with the highest year at $34 billion. While Scenario 3 applied the same startup fixed capital investment (FCI) subsidies to all infrastructure-compatible fuel pathways, Scenario 4 saves the most favorable startup FCI subsidies for fast pyrolysis. Because these FCI subsidies end up being directed toward fast pyrolysis, the subsidies are more efficient in promoting take off of that pathway than if subsidies are spread to different pathways or if subsidies are directed to a pathway that is not as economically attractive. As a result, even though a $10 billion annual subsidy was allowed, subsidies in Scenario 4 never surpassed $6 billion in any year after 2012.

Fast pyrolysis receives FCI subsidies for pioneer and commercial plants in 2012-14. After that time, cellulosic ethanol becomes a more attractive investment because it starts with more learning at the pilot and demonstration levels and builds on that head start. Only when this initial wave of fast pyrolysis plants is built and generates its own learning do fast pyrolysis plants again become attractive investments. Additional fast pyrolysis commercial plants are built after 2017 without a subsidy. The initial subsidies in 2012-14 are enough to set in motion the learning necessary to make commercial plants an attractive investment without additional FCI subsidies. The threshold volume of 0.3 billion gallons for the FCI subsidy for commercial plants is reached in 2017. As a result, after 2014, fast pyrolysis is supported by only the price subsidy and by loan guarantees.

Although additional money was available to spend in each year, greater spending on subsidies did not result in substantially more output. Spending more money on subsidies toward fast pyrolysis ends up subsidizing plants that would have come online without subsidies and/or result in more output but at a rate of spending above $10 billion per year. Additional subsidies available to pathways other than fast pyrolysis have little effect. By restricting subsidies to fast pyrolysis, it becomes more mature and locks out other pathways.

Though it appears most economically efficient, relying on this single pathway presents nontrivial technology risks. Relying on a pathway with unfavorable long-term economics could result in less volumetric output. In the BSM, the most economically attractive pathway is obvious based on the available input data. In reality, the consequences of choosing a less than ideal pathway may not be evident until several years after a policy decision.

Insight 4: Technologies with favorable long-term economic cost structures can succeed if supported by targeted subsidies.

BSM simulations have shown that technological "lock-in" is likely to occur. Fischer-Tropsch has the highest initial level of maturity among the infrastructure-compatible fuels; its initial settings for pilot-scale and demo-scale maturity are higher than or equal to all other pathways (with the exception of the starch and cellulosic ethanol pathways). However, the mature commercial plant economics of fast pyrolysis are better than that of Fischer-Tropsch, based on the available process designs. In Scenario 4, the maturity of fast pyrolysis has to increase in order to make the investment look more attractive and to prevent technology lock out from Fischer-Tropsch.

In Scenario 5, subsidy policies are crafted to avoid lock out by any one pathway. To overcome lock out, subsidies target learning through pioneer plants and do not include commercial plants. By staggering policy start times and varying durations of subsidies according to maturity, the other technologies have a chance to build experience. Limiting more mature or economically attractive technology subsidies to begin after the volumetric threshold is reached allows the other technologies to also develop. Rather than pouring extra subsidies into a relatively mature technology, this approach provides the minimum subsidies needed for a more mature technology (such as Fischer-Tropsch) to develop on a commercial scale, while providing the others with the extra support they need to accelerate their experience levels. This approach allows the successful take-off of four technologies while also approaching RFS2 production levels in 2030.

However, though heavy subsidies may help overcome initial maturity differences, they are not necessarily sufficient in overcoming differences in long-term economic cost structures. Technology, such as fermentation (which is as commercially mature as the other pathways by the end of the simulation), may need support beyond the subsidies exercised in this analysis order to reach greater production levels.

Current and potential use of BSM

The BSM has provided an invaluable tool for the Bioenergy Technology Office of the DOE for gaining intuition around the biomass-to-biofuels supply chain, and the insights detailed herein only begin to address the impact the project has had in building understanding around several key industry dynamics. The BSM has also been utilized in collaboration with other parties, such as the EPA. Although to date the model has been mainly used by the DOE and other governmental agencies, it has the potential to be highly useful to many different stakeholders and across a wide range of analysis areas within the biofuels industry, as shown in Table 6.

Table 6. Potential collaborations for ongoing BSM use

Stakeholder / Analysis Area	**Climate Change**	**Supply Curves**	**Biomass Yield**	**Policy**	**R&D**	**Energy Security**	**Trade**	**Region-specific**
EPA	X			X				
USDA	X	X	X	X				
DOD						X		
Oil Companies				X	X			
Biofuels Companies				X	X			
Think Tanks	X	X	X	X	X	X	X	X
Foreign Governments/ Organizations							X	X
Universities	X	X	X	X	X	X	X	X

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 infrastructure-compatible 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 an operational structure that reflects both the physics and economics of the system, BSM is a tool for building understanding around initiatives that seek to stimulate sustained development of the industry. And by representing the system of interactions simply and transparently, the model sheds light on gaps in the data as well as areas where understanding of system structure is in need of enrichment. Analyses of the BSM—both as standalone modules and in integrated form—have underwritten powerful insights about the nature of the biomass-to-biofuels supply chain and of the nature of policy initiatives required to stimulate industry take off.

Over the course of this project, the BSM working team has developed numerous internal reports and briefing documents, which are housed on the project repository at (bsm.nrel.gov). The References section of this paper also highlights several publications based on the BSM project.

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

Price mechanisms within BSM can be viewed 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 31.

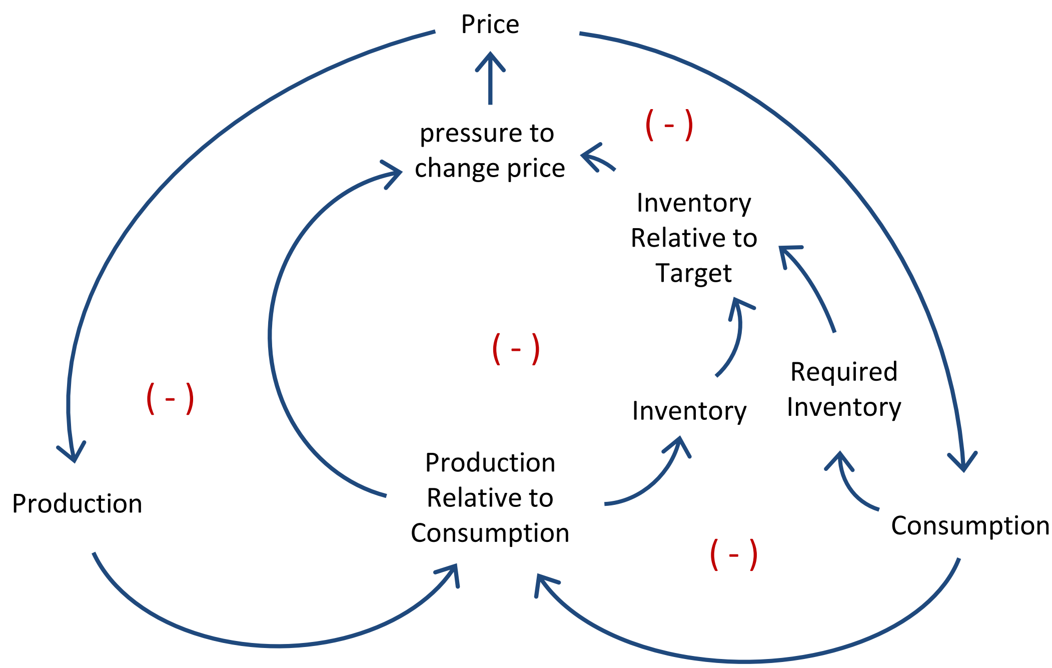

Figure 31. Stylized view of feedbacks in BSM pricing mechanism

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 32 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.

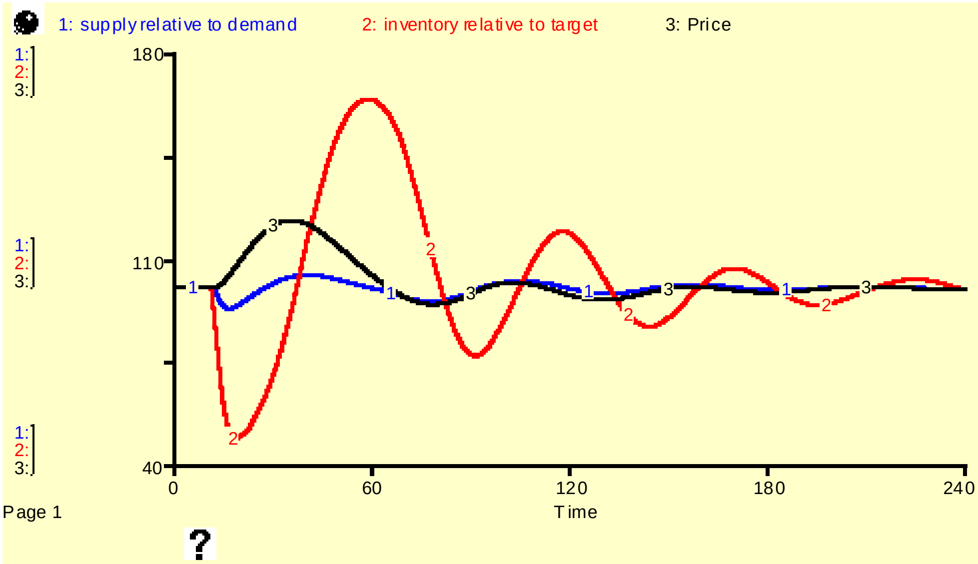

Figure 32. Response of pricing system to 10% step-increase in demand

In BSM pricing, the mechanisms that connect production, consumption, and inventory to fractional change in price are significantly more detailed. The structural arrangement shown in Figure 33 is used to determine dynamic prices of several products throughout BSM, including:

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

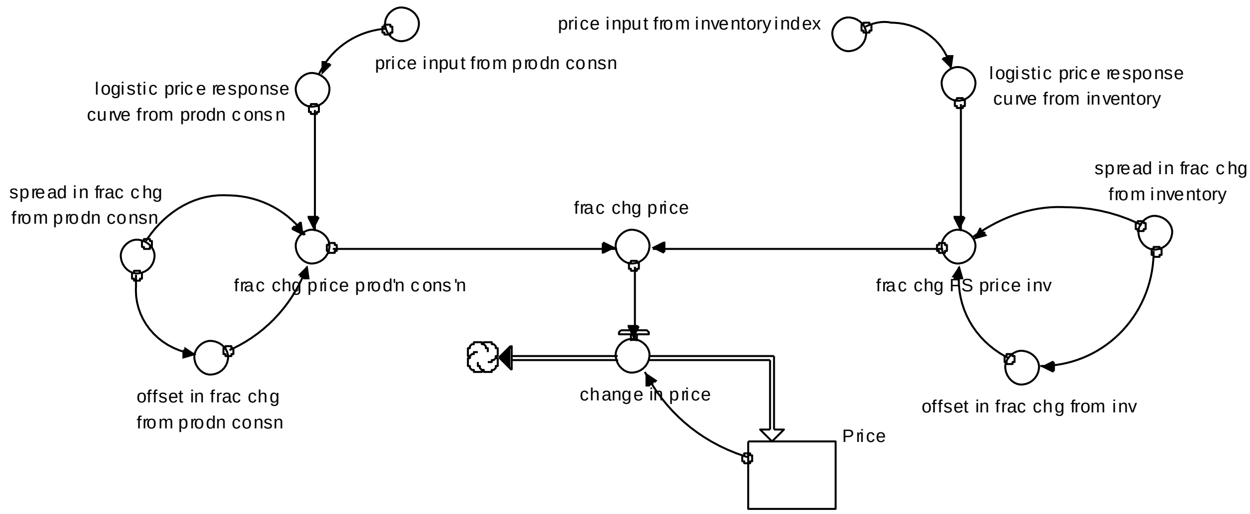

Figure 33. Detail of BSM generic pricing structure

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 or 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 33.

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 34.

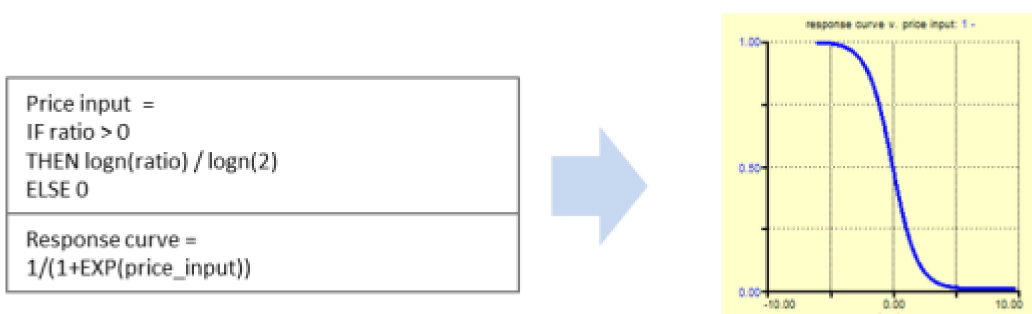

Figure 34. Illustrative price input and response curve calculations

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, with context-specific details (beyond the scope of this paper) applying to specific fuel markets.

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:

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

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:

- It 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, as summarized in Table 7.

Table 7. 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 ethanol by butanol in lo-blend mixes Displacement of non-bio-butanol in industrial market	Butanol, ethanol prices	
Fuel Use	Allocation of fuel sales between hi-blend and gasoline	Price of gasoline Price of high-blend	

Appendix C: Aggregation of age classes in the vehicle module

The current 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 the current version of the model, we have reduced computational overhead significantly by aggregating age distribution of vehicles. We represent vehicle vintages using 5 distinct cohorts. Each cohort represents the 4 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 35.

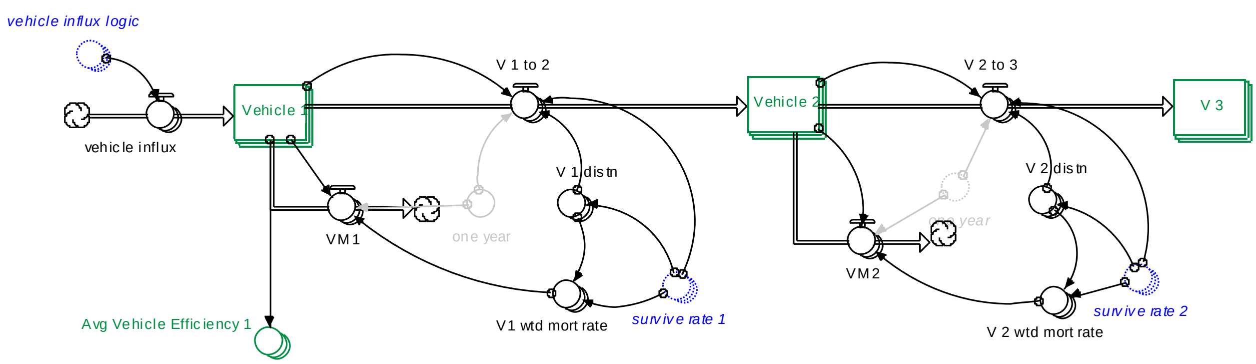

Figure 35. Structure of vehicle vintaging

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

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

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

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

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

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

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

Figure 36 compares the transient response of a single 4 year cohort of the BSM vintaging structure against a simple one-stock structure and against a more disaggregated structure with 4 1-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.

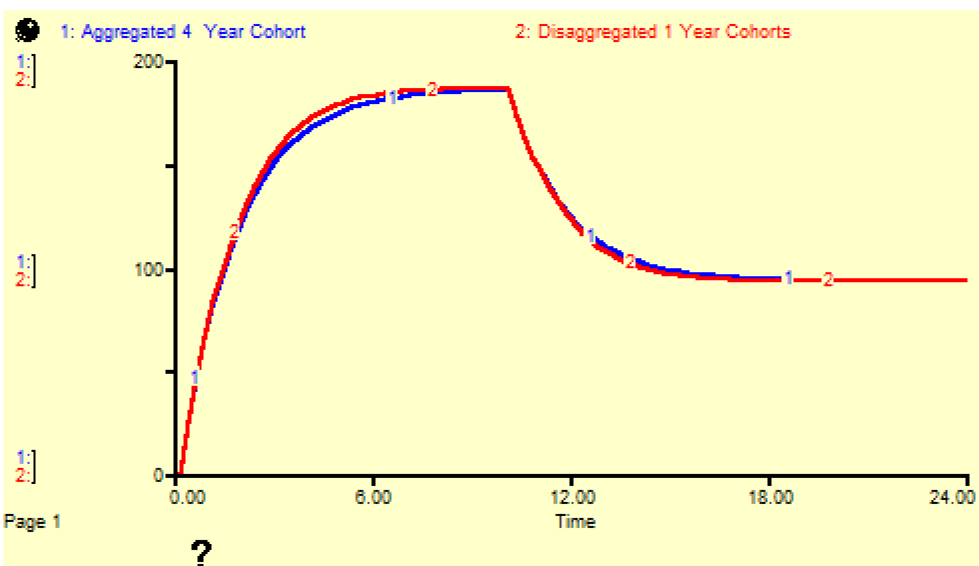

Figure 36. Comparison of BSM and disaggregated vehicle cohorts.

References

Anderson, S. T. (August 2009). *The Demand for Ethanol as a Gasoline Substitute*, Michigan State University, East Lansing, MI

Bush, B., M. Duffy, D. Sandor & Peterson, S. (2008). *Using System Dynamics to Model the Transition to Biofuels in the United States: Preprint*, National Renewable Energy Laboratory, Golden, CO

Energy Information Administration. (May 2010). Annual Energy Outlook 2010, In: *Forecasts & Analyses,* 22.02.2011, Available from http://www.eia.doe.gov/oiaf/archive/aeo10/index.html

Hax, A. C., & Majluf, N. S. (1982). Competitive Cost Dynamics: The Experience Curve. *Interfaces*, 50-61.

Interagency Agricultural Projections Committee. (2007). *USDA Agricultural Projections to 2016*. U.S. Department of Agriculture

isee systems. (2010). STELLA: Systems Thinking for Education and Research Software, 22.02.2011, Available from
http://www.iseesystems.com/softwares/Education/StellaSoftware.aspx

Johnson, C., & Melendez, M. (2007 draft). *E85 Retail Business Case.* National Renewable Energy Laboratory.

Johnson, C., & Melendez, M. (2007 draft). *E85 Retail Business Case.* National Renewable Energy Laboratory.

Johnson, C., & Melendez, M. (2007 draft). *E85 Retail Business Case.* National Renewable Energy Laboratory.

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

Merrow, E. W. (1983). *Cost Growth in New Process Facilities.* Rand Corporation.

National Association of Convenience Stores. (2007). *Annual Report.*

Newes, E., Inman, D., & Bush, B. Understanding the Developing Cellulosic Biofuels Industry through Dynamic Modeling. In M. e. dos Santos Bernardes, *Economic Effects of Biofuel Production.* InTech Open Access.

Office of the Biomass Program and Energy Efficiency and Renewable Energy. (March 2008). *Multi-Year Program Plan March 2008,* U.S. Department of Energy, Washington, D.C.

RW Beck. (2010). *Biorefinery Learning Curve Analysis.* Report for NREL Subcontract LCI-9-88660-01.

Sterman, J. (2000). *Business Dynamics.* Irwin/McGraw-Hill.

United States Department of Agriculture. (n.d.). *Agricultural Baseline Projections*. Retrieved from ERS/USDA Briefing Room: http://www.ers.usda.gov/Briefing/Baseline/

United States Energy Information Agency. (n.d.). *EIA Analysis and Projections*. Retrieved from http://205.254.135.24/analysis/projection-data.cfm#annualproj

The University of Tennessee. (n.d.). Agricultural Policy Analysis Center Research Tools - POLYSYS, In: *Agricultural Policy Analysis Center - The University of Tennessee,* 22.02.2011, Available from http://www.agpolicy.org/polysys.html

U.S. Department of Agriculture. (February 2011). Conservation Reserve Program, In: *Conservation Programs,* 22.02.2011, Available from http://www.fsa.usda.gov/FSA/webapp?area=home&subject=copr&topic=crp

U.S. Environmental Protection Agency. (December 2010). Renewable Fuels: Regulations & Standards, In: *Fuels and Fuel Additives,* 22.02.2011, Available from http://www.epa.gov/otaq/fuels/renewablefuels/regulations.htm

U.S. Government. (2007). *Energy Independence and Security Act of 2007*

Vimmerstedt, L., Bush, B., & Peterson, S. (2012). Ethanol Distribution, Dispensing, and Use: Analysis of a Portion of the Biomass-to-Biofuels Supply Chain Using System Dynamics. *PLOS ONE.*