Table of Contents

Go Back

Simulating the Impact of a Strategic Fuels Reserve in California

Professor Andrew Ford
Program in Environmental Science and Regional Planning
Washington State University
Pullman, WA 99164-4430

Submitted to the
21st International Conference of the System Dynamics Society

May 5, 2003

Preface

In March of 2003, armed forces from the US and the UK went to war in Iraq. Crude oil prices had climbed to over $37 per barrel early in the month and had fallen to around $27 per barrel by March 24th. By mid April, crude oil was selling for under $25 per barrel. As these events unfolded, California officials were examining unusually high gasoline prices which had “soared in advance of the war in Iraq” (Los Angeles Times, March 24, 2003). Sen. Barbara Boxer asked the General Accounting Office to probe the gasoline situation. Rep. Bob Filner introduced legislation to freeze gasoline prices nationwide and launch an investigation by the Secretary of Energy. Some officials believe the “time is right to push for measures that may shield motorists from soaring costs.” Key agencies in California are said to be preparing to “lay out a package of recommendations” by the end of April and to present a “final strategy” to the Governor and the Legislature in June. One of the “most ambitious options” would create a “gasoline bank” that would act as a reserve supply when refineries falter.

A California fuels reserve is not a new idea. The Los Angeles Times reports that the idea dates to the “last gasoline squeeze in 1999, when refinery outages sent California gas prices soaring.” After an investigation, Attorney General Bill Lockyer declared California’s oil and gasoline markets “deeply flawed, hampered by anemic competition and unique characteristics that preclude outside aid when problems arise” (Los Angeles Times, March 24, 2003). “Contending that government has a role to play when market forces fail, Lockyer pressed for studies, including one exploring whether a strategic fuels reserve could help stabilize the market.” Key studies were completed early in 2002, but they are reported to have “received scant attention because by then gas prices were back to normal levels.” But the recent price run-up has given new relevance to the proposals, and the Los Angeles Times reports that lawmakers in both Washington DC and in Sacramento “seem primed to consider bold moves.”

Background on AB 2076

Gasoline prices in California are more volatile than in the rest of the country due to a variety of factors including the insular nature of the California market and the unique fuel specifications to improve air quality (Stillwater 2002, p 5). According to Finizza (2002, p 1), volatility in gas prices has increased since the introduction of CARB Phase II gasoline and has remained at high levels since 1999. In 1999, following a series of refinery outages and price spikes, the California Attorney General created a task force to investigate price volatility. The efforts of the task force led to a report on gasoline pricing in California (AG 2000) and to Assembly Bill 2076.

AB 2076 called on the California Energy Commission (CEC) to examine the feasibility of operating a strategic fuel reserve (SFR). The CEC was directed to study the costs and benefits of a SFR designed to insulate California consumers and businesses from short-term price increases arising from refinery outages and other similar supply interruptions. AB 2076 also directed the CEC to examine an appropriate level of reserves, adding "in no event may the reserve be less than the amount of refined fuel that the commission estimates could be produced by the largest California refinery over a two-week period" (Stillwater 2002, p xv). The CEC issued two reports on the SFR in July of 2002 (Stillwater 2002, Finizza 2002).

The Stillwater report concludes that "the California gasoline market suffers from insularity caused by its unique specifications, a subsequent lack of liquidity, inability to lock in pricing for forward trades, and impediments to market entry by outside sources" (Stillwater 2002, p 1). These factors are listed along side of the "supply disruptions identified as a cause of price spikes in the legislation" that led to the Stillwater study. The Stillwater report found "overwhelming evidence"of the consumer benefits associated with the creation of an SFR (Stillwater 2002, p 140). It called for the state of California to "issue a tender for the creation of a 5 million barrel of versatile petroleum product storage under long-term lease agreements" (Stillwater 2002, p 3). A central principle is a SFR that operates free of government interference with market forces:

> *Unlike European, Asian and US federal reserve systems, California SFR inventories will not be sold at all. Its gasoline will be 'time-traded.' A barrel out equals a barrel in. Contractual volumes will be loaned out on the next pipeline cycle for replenishment within 'x' number of weeks. The SFR will be a rolling inventory that provides a physical basis for greater forward liquidity …The loan out of SFR gasoline supplies will serve to bridge the time gap to other markets.*
>
> (Stillwater 2002, p B-4)

In the second report, Finizza describes a mathematical model of the price impacts of refinery disruptions in California. The model allows for refinery disruptions of different size and duration, and it calculates the price impact on California consumers. The SFR is represented by a limit on the price impact based on an estimate of the "time swap auction premium,” the portion of the price spike that an SFR could eliminate. The model calculates expected impacts across a wide range of disruptions with the help of “Crystal Ball,” a Monte Carlo estimator add-in to Excel. Finizza found “the net benefit of the SFR to the California consumer of avoiding prices spikes” at around “$400 million per year against an annualized cost of $20 million.” He estimated that the benefits can “rise to $700 million or fall to below $200 million” and noted “even at the low value, the benefits are an order of magnitude above the projected costs.”

Purpose of this Paper

This paper describes a system dynamics model to simulate the changes in the spot price of gasoline following the disruption of a refinery in California. The model was designed as a dynamic expansion of the spreadsheet model by Finizza. The system dynamics model provides a more detailed description of the SFR, and it allows the CEC to analyze a variety of other options for mitigating the impact of refinery disruptions. The model is based on discussions with staff at the California Energy Commission (CEC). It draws on information in the reports by Stillwater and Finizza and the theory of commodity spot markets published by Pindyck (2001).

Introduction to the Model

Figure 1 shows the opening screen of the model interface. The interface aids the user in the selection of inputs and in the display and understanding of results. The opening screen contains four navigation buttons to allow the user to jump to different screens for operating the model. The opening screen introduces the user to the purpose and focus of model. The large button (with the photo of gasoline storage tanks) explains that storage is at the center of the model. The draw down and replenishment of private storage are simulated inside the model. The model assumes that most of the gasoline in storage serves as operational reserves. The remainder serves as precautionary (or strategic) reserves.

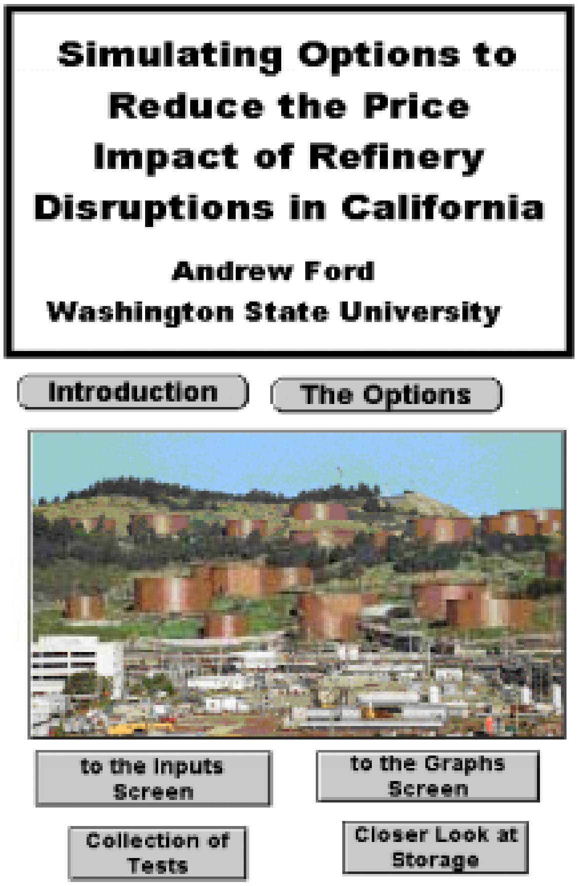

Simulating Options to Reduce the Price Impact of Refinery Disruptions in California

Andrew Ford
Washington State University

Figure 1. Opening Screen of the Interface.

The model also simulates the gasoline stored in the SFR based on the Stillwater proposal for a time swap mechanism to provide prompt supplies of gasoline during the aftermath of a disruption. (Policy makers are interested in a variety of options for mitigating the price impacts of refinery disruptions including. Because of space limitations, this paper focuses on the SFR.)

The model is designed to promote multiple experiments with different values of the input parameters shown in Figure 2. This screen provides control over the main inputs, and it provides information buttons to aid in the selection of the inputs. Figure 2 shows that the refinery capacity is set at 1,000 TBD (thousand barrels per day of gasoline). The dials have been set for a 15-day disruption of 100 TBD of refinery capacity starting in the 20th day of the simulation. The system would lose 10% of the refinery capacity and around 1.5 million barrels of production. The simulation begins with retail demand at 1,180 TBD, the "base case" estimate for 2002 by Stillwater (2002, p. 20). The demand responds to changes in retail prices based on a price elasticity of –0.1, the value used by Finizza (2002, p 74). The controls in Figure 2 call for an 18-day lag for passing wholesale price increases to the retail level. The lag for passing wholesale price reductions to retail is set at 28 days. These lags were selected to provide a good match with the asymmetric price response described by Finizza (2002, p 34).

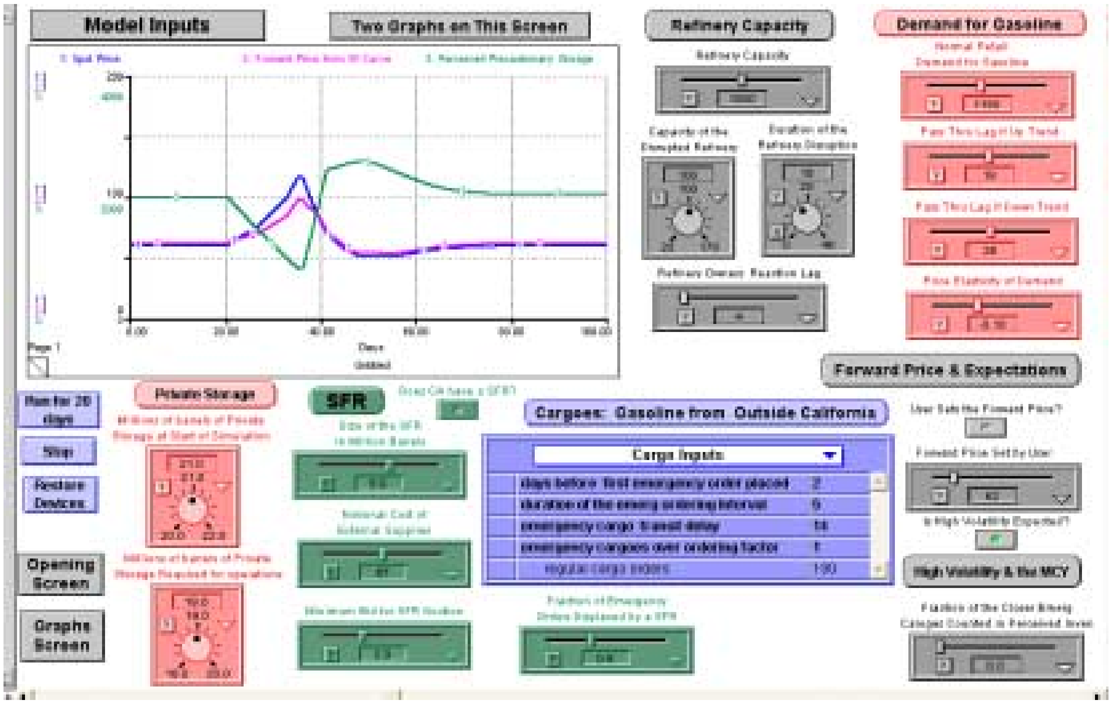

days before first emergency order placed	2
duration of the ensuing ordering interval	5
emergency cargo transit delay	14
emergency cargoes over ordering factor	1
regular cargo orders	100

Figure 2. The inputs screen at the conclusion of a simulation of a 15-day disruption of 100 TBD of refinery capacity.

The upper dial in the private storage portion of Figure 2 sets the initial value of gasoline in private storage at 21 million barrels. The lower dial sets the operating reserve requirement at 19 million barrels. These dial settings ensure that the simulation begins with 2 million barrels of "discretionary reserves," a plausible starting point located midway in the range of estimates by Stillwater (2002, p 65). Other inputs in Figure 2 call for 190 TBD of regular cargoes which arrive during each day of the simulation. With 190 TBD of cargoes from outside California and 1,000 TBD of refinery capacity, the system could supply 1,190 TBD of gasoline. With the retail demand set at 1,180 TBD, the system would need to operate refinery capacity at 99% to satisfy the retail demand. The model distinguishes between "regular cargoes" and "emergency cargoes." Regular cargoes are planned well in advance of any disruption and arrive on a continuous basis during the 100-day simulation. The emergency cargoes are ordered in the days following a disruption to make up for the estimated lack of production. The inputs in Figure 2 show that the orders begin 2 days after the disruption; are spread over a 5-day interval; and cargoes arrive after a 14-day transit delay.

An Illustrative Simulation

Figure 3 shows the "graphs screen" of the model for the simulation with a 100 TBD disruption. For clarity, we assume that the system is in equilibrium during the first 20 days of the simulation. The spot price increases rapidly during the 15-day disruption. It peaks in the 35th day and descends to values below the initial equilibrium. The simulation concludes with a slow recovery to the initial equilibrium conditions.

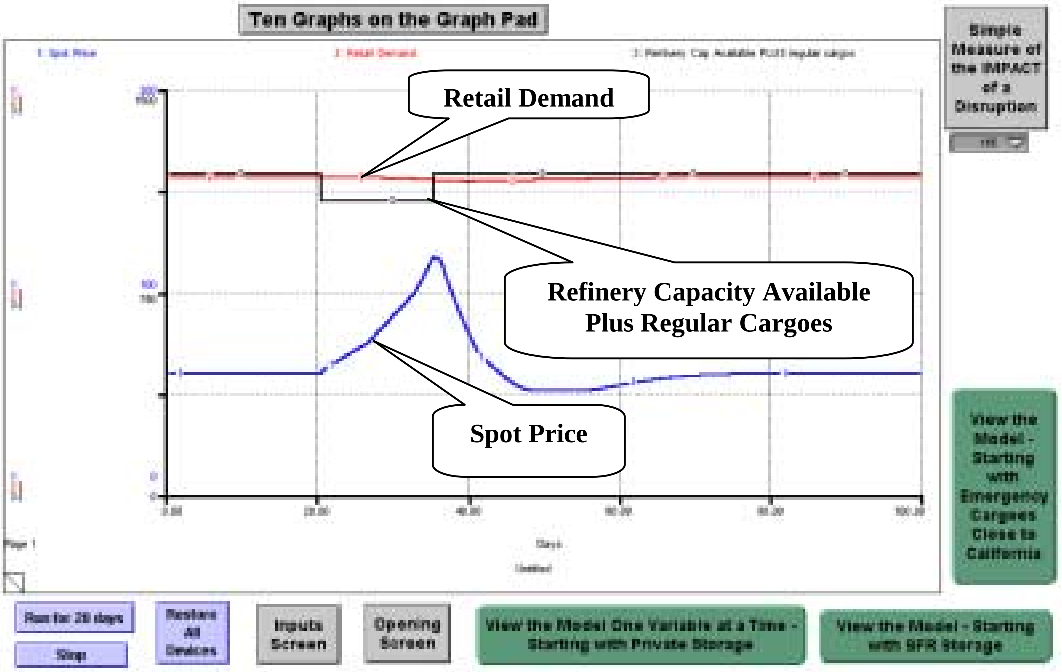

Figure 3. The graphs screen of the model interface with the first of ten graphs in view.

The numerical display in the upper right corner of Figure 3 shows the impact of the simulated disruption at $165 million. The model keeps track of wholesale payments as if 40 million gallons per day were sold at the spot price over the entire simulation (Finizza 2002, p. 66). If there were no disruption, the payments are $2,400 million. In the disruption shown in Figure 3 the cumulative payments are $2,565 million, $165 million more than if there were no disruption.

Figure 4 shows the second of ten graphs on the graphs screen. The new graph shows the simulated behavior of "actual precautionary storage" on a scale from 0 to 4,000 TB (thousand barrels). Precautionary storage is drawn down during the 15-day interval when the 100 TBD of refinery capacity is out of service. The loss of production is 100 TBD times 15 days or 1.5 million barrels. (Some call the 1.5 million barrels the "disrupted barrels.") Figure 4 shows that the precautionary storage is reduced during the 15 days, but not by the entire 1.5 million barrels that one might expect from the "disrupted barrels." The decline in storage is mitigated somewhat by a small increase in refinery utilization and a small decline in retail demand.

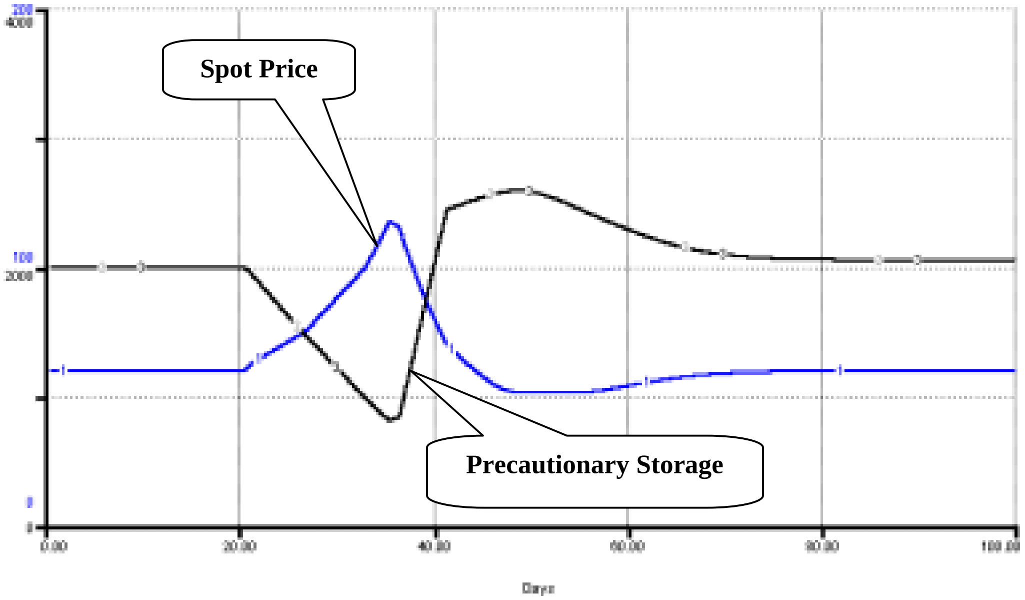

Figure 4. Illustrative results for the spot price and the precautionary storage.

Figure 4 shows that precautionary storage increases rapidly around the 35th day of the simulation, the time when the capacity is back in operation and emergency cargoes arrive in California. Precautionary storage reaches around 2.6 million barrels by the 50th day of the simulation. This is the period when spot prices are at their minimum value. The low spot prices cause refinery utilization to fall below 99% allowing the precautionary storage to gradually return to the 2 million barrels seen at the start of the simulation.

Figure 5 shows the retail price compared with the spot price with the vertical scaled from 0 to 200 cpg, the same scale used previously. The model assumes that the retailers have 90 cpg in fixed expenses, so the simulation begins with the spot price at 60 cpg and the retail price at 150 cpg. The spot price increases rapidly during the 15 days of the disruption, peaking at 117 cpg in the 35th day of the simulation. There is an 18-day lag in passing the wholesale price increases to the retail level. The retail consumer would experience the greatest impact around the 40th day. The retail price has risen by around 13% by this time, and the consumers react with a 1.3% reduction in gasoline consumption. This reaction is too little and too late to make a difference in the upward trajectory of the spot price in Figure 5. The small reduction in retail demand appears in the simulation around the time that emergency cargoes have arrived, the refinery is back in operation and the spot price is descending rapidly. With these assumptions, the demand-side response will lead to a somewhat lower prices when the system “bottoms out” around the 45th day of the simulation.

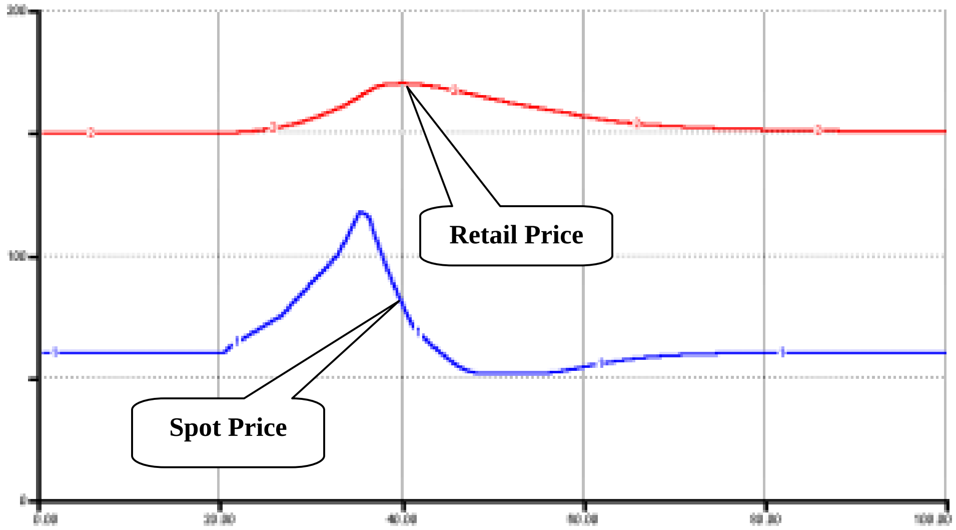

Figure 5. Spot price and retail price during the illustrative simulation.

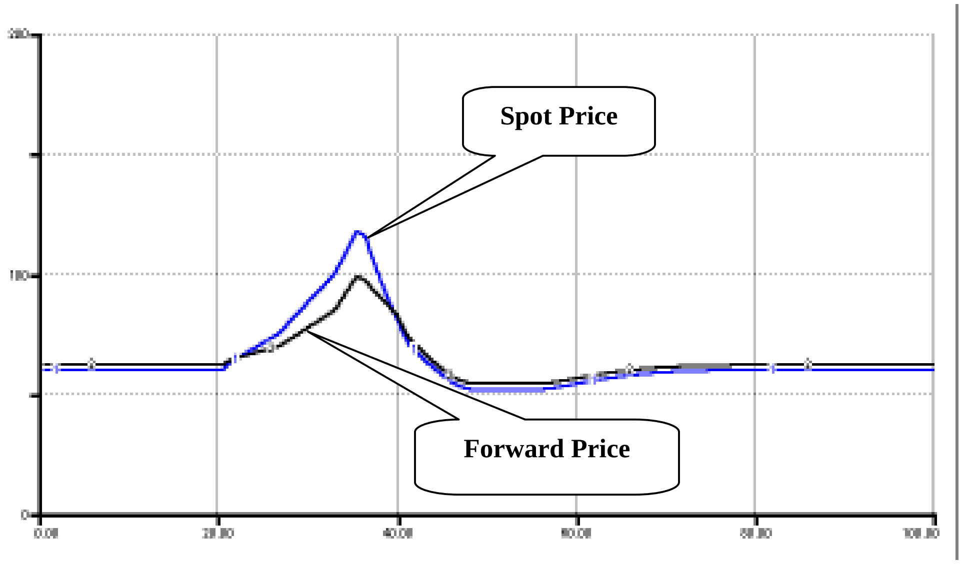

Figure 6. Spot price and forward price during the illustrative simulation.

Figure 6 shows the forward price compared with the spot price. The simulation begins with the spot price at 60 cpg and the forward price at 62 cpg. The forward price refers to delivery of gasoline one month in the future. The forward price exceeds the spot price by 2 cpg; the markets may be said to show a “spread” of 2 cpg. The forward price is calculated as an endogenous variable based on the spread to be expected from a "Working Curve" (Working 1949). With 2,000 TB of storage, the expected spread is 2 cpg, a situation called “contango” in the language of the market (Duffie 1989, p 101). According to Keynes (1930, Vol II Ch 29) and Duffie (1989, p 101), a contango is to be expected when there are surplus stocks and “this contango must be equal to the cost of the warehouse, depreciation and interest charges of carrying the stocks.”

Figure 6 shows that the market situation changes shortly after the refinery outage appears in the 20th day of the simulation. Within a few days, the spot price has increased sufficiently to exceed the forward price. According to Pindyck (2000, p 17), the markets may be said to exhibit "strong backwardation." By the 30th day, for example, the spot price is at 90 cpg, and the forward price is at 78 cpg. Figure 6 shows that the two prices peak around the 35th day of the simulation. The spot price reaches 117 cpg while the forward price reaches 99 cpg. Within a few days, however, the spot price has descended rapidly, and the markets have returned to the contango situation seen at the start of the simulation.

Stocks and Flows of Gasoline

Figure 7 shows the stock variable to represents the "Gasoline in Private Storage." The stock is measured in TB, thousands of barrels. The stock is influenced by two flows, each measured in TBD, thousands of barrels per day. The stock is increased by the flow "excess production adds to storage" and it is reduced by the flow "withdrawals from storage." The double lines in Figure 7 represent the flow into or out of the stock. In this case, the flows are easy to visualize because they represent the flow of gasoline.

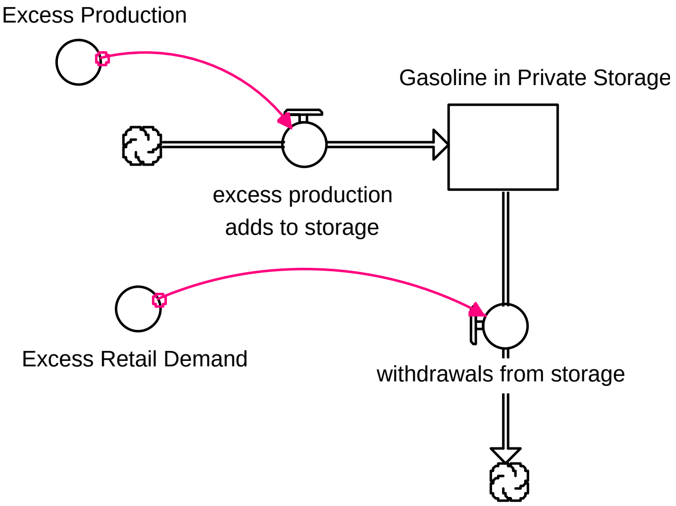

Figure 7. Stock of Gasoline in Private Storage.

The remaining variables in Figure 7 are called converters. Converters are used to help explain the flows. In this example, the flow feeding into the stock is identical to the "Excess Production." (In this case, we are introducing an additional variable with a longer name to make the meaning clear.) The flow draining the stock is influenced by the "Excess Retail Demand."

Figure 8 shows three stocks to represent gasoline in transit toward California. These stocks are designed with vertical slats to remind us that they are "conveyors", a special category of stock whose outflow is controlled by the timing of the inflow and the length of the transit time (Ford 1999, Chapter 10). The model distinguishes between two types of cargoes arriving from outside of California. The regular cargoes arrive at a constant rate regardless of the disruption. The lower part of Figure 8 shows two stocks to represent the emergency cargoes on the way to California. Orders for emergency cargoes would be placed shortly after the disruption, and there is a 14-day transit delay before they are received in California and contribute to the "supply from cargos."

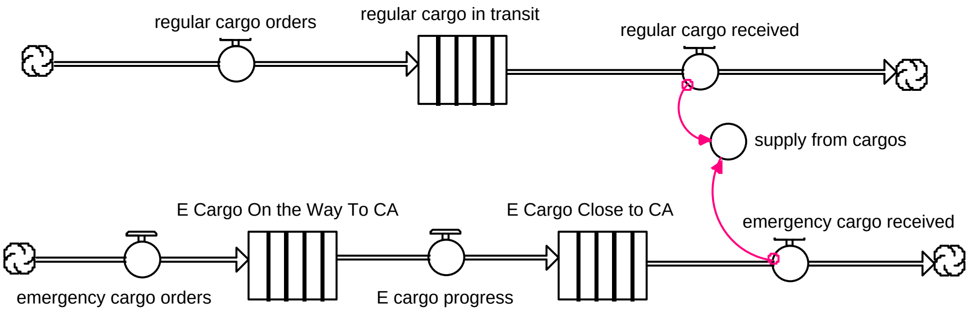

Figure 8. Stocks of gasoline in transit toward California.

Figure 9 shows the stocks and flows to simulate the SFR. The stock of Gasoline in the SFR is reduced when gasoline is released because of requests for prompt supply. Replacement orders are issued simultaneously and build the stock of "Gasoline on the Water." After a "Time Swap Return Interval," the replacements are received at the SFR.

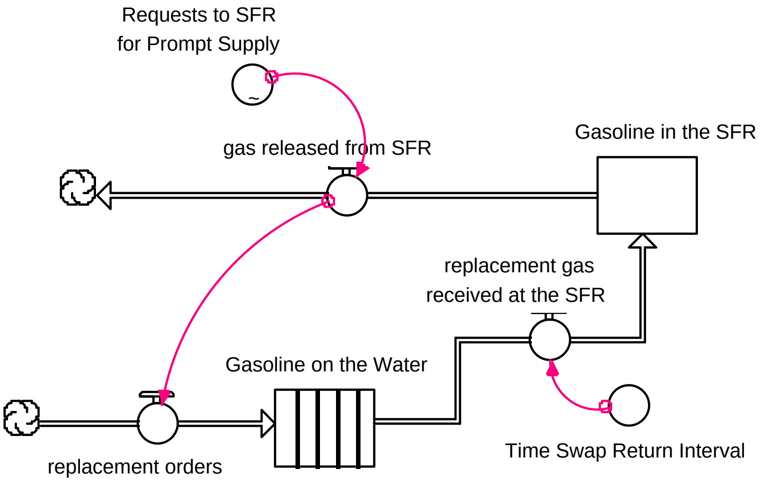

Figure 9. Stock for Gasoline in the SFR.

Simulating the Spot Price Response to Supply and Demand

The stocks and flows in the previous three diagrams are easy to visualize because it is easy to visualize the flow of gasoline from one location to another. In Figure 10, we see a stock variable used in a somewhat different manner. In this case, the stock represents the spot price measured in cpg. The flow is the "spot price change" which is measured in cpg/day. This flow is a "bi-flow;" it can cause the spot price to rise or to decline. The flow is controlled by the "fractional daily change in price." If the supply falls short of the demand, the fractional change is positive, and the spot price increases. If the supply exceeds the demand, the fractional change is negative, and the spot price declines. The "Supply of Gasoline" is measured in TBD and is the sum of

- refinery production,
- supply from cargoes,
- withdrawals from private storage, and
- supply from the SFR.

The "Demand for Gasoline" is the sum of the retail demand and the wholesale demand to rebuild inventory.

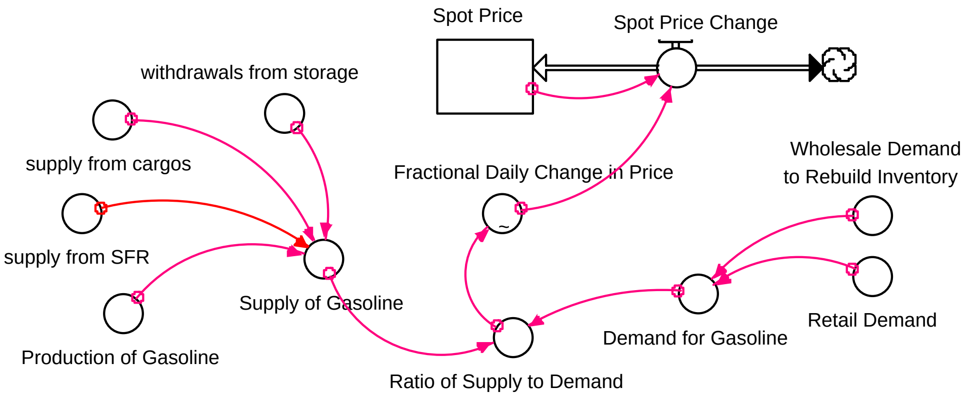

Figure 10. Stock variable used to represent the spot price of gasoline.

Simulating the Draw Down of Private Storage

Figure 11 shows some of the main variables to help the reader understand the theory of storage draw down implemented in the model. The stock of gasoline in private storage is 21,000 TB at the start of the simulation. The actual gasoline in storage is added to the “Emergency Cargoes Counted in Perceived Inventory” to give the “Perceived Storage.” For now, let’s assume that traders are not reacting to emergency cargoes as if they are already in California. With this assumption the Perceived Storage is the same as actual storage, 21,000 TB at the start of the simulation. With 19 million barrels needed for operating reserves, the Perceived Precautionary Storage is 2,000 TB.

The Perceived Precautionary Storage is used to obtain two estimates of the MCY, the marginal convenience yield measured in cpg/month. According to Pindyck (2001, p 6), the MCY is the value of the flow of services accruing from holding the marginal unit of inventory. Pindyck uses the terms “marginal convenience yield,” "price of storage" and “marginal value of storage” synonymously. He states that the price of storage will rise sharply as the private storage falls:

> *The marginal value of storage is likely to be small when the total stock of inventories is large (because one more unit of inventory will be of little extra benefit), but it can rise sharply when the stock becomes very small. Thus we would expect the demand for storage function to be downward sloping and convex.*
>
> (Pindyck 2001, p 7)

Pindyck shows two downward sloping, convex curves and advises that the higher of the two curves would apply if the spot market exhibits high volatility.

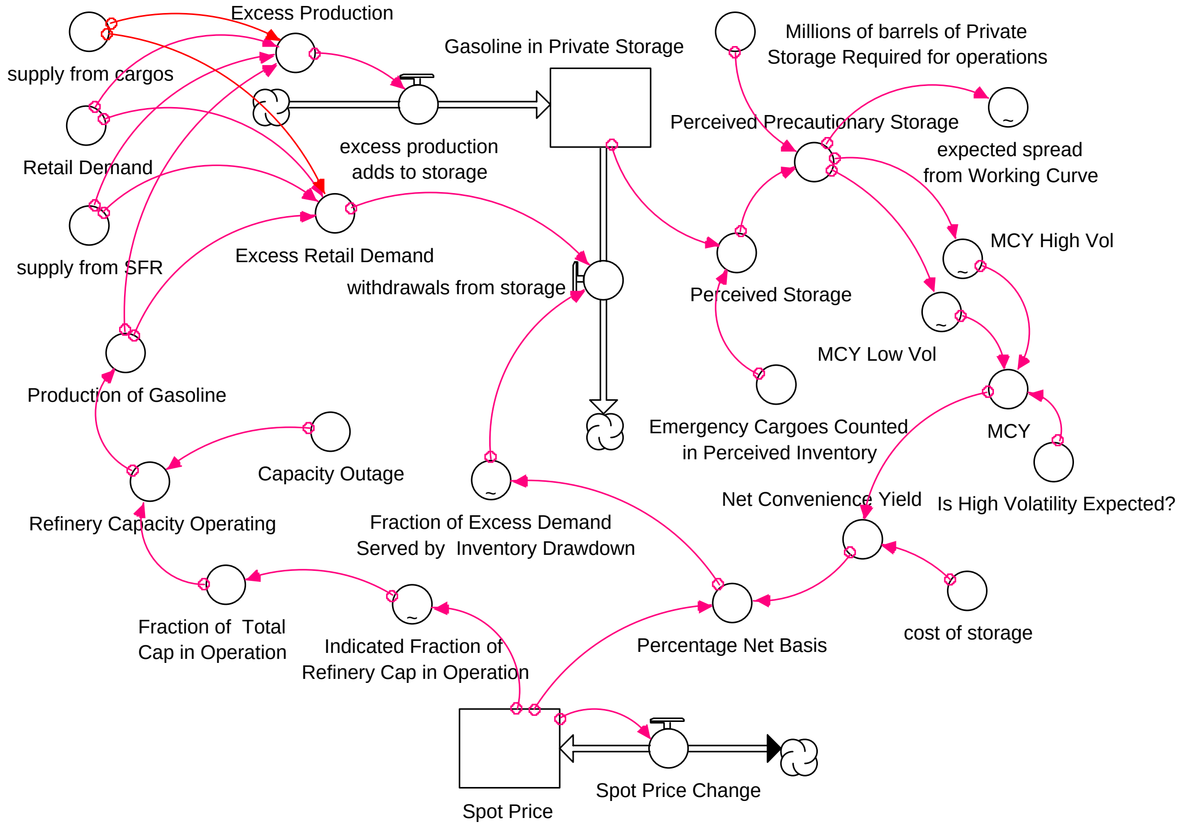

Figure 11. Model variables controlling the draw down of private storage.

Pindyck gives only general advice on the convex shape of the MCY curve. He argues that one can estimate the net convenience yield using an equation that must be satisfied if the markets are to “avoid arbitrage opportunities.” Pindyck’s equation (6) gives the marginal convenience yield over a specified period of time based on the futures prices, risk-free interest rate, spot prices, and unit cost of physical storage. He presents estimates of the convenience yield for gasoline over the interval from 1984 to 2000. A typical value of the inferred spot price is 60 cpg; a typical value of the convenience yield is 5 cpg/month. Pindyck (2001, p 22) describes the convenience yield for gasoline as “economically significant.” His measure of significance is to compare the convenience yield to the spot price. The monthly marginal convenience yield is said to be 8.1% of the mean spot price.

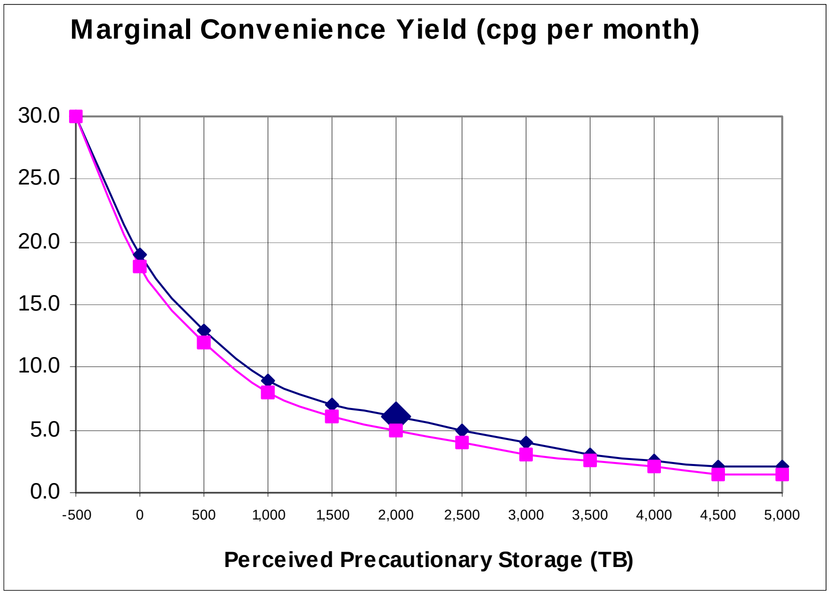

Figure 12. Nonlinear Curves for the MCY.

The large diamond on the upper curve in Figure 12 draws our eye to the starting conditions: the MCY is 6 cpg/month when the perceived precautionary storage is at the standard value of 2,000 TB. With the cost of storage at 1 cpg/month, the starting value of the net convenience yield is 5 cpg/month, the mean value estimated by Pindyck. As storage is drawn down in the days following a disruption, the MCY "moves up the curve." Later in the simulation, storage will increase rapidly (with the arrival of emergency cargoes), and the MCY will move rapidly down the curve.

Figure 11 shows that the MCY is used to determine if gasoline will be removed from private storage during a disruption. The model assumes that 8 %/month may be used as a standard value of this key indicator. If owners of storage see a Percentage Net Basis of 8 %/month, we assume they are willing to release gasoline into the spot market. But if the Percentage Net Basis rises above 8 %/month, gasoline is more valuable held in storage. Owners would be reluctant to release gasoline into the spot market under such conditions. This approach is based on Pindyck's interpretation:

> *The ratio of the net convenience yield to the spot price is referred to as the percentage net basis, and is analogous to the dividend yield on a stock. In fact, if storage is always positive, one can view the spot price of a commodity as the present value of the expected future flow of convenience yield, just as the price of a stock can be viewed as the present value of the expected future flow of dividends.*
>
> (Pindyck 2001, pages 16-17)

This behavioral assumption is implemented with the shape of the nonlinear curve for the "Fraction of Excess Demand Served by Inventory Drawdown." This variable is connected to the flow "withdrawals from storage" which, in turn, influences that amount of gasoline in private storage in the future.

Information Feedback in the Model

Figure 13 shows the two loops that the control the main supply and demand for gasoline in the simulation model. The TBD of retail demand and refinery production are much larger than other variables in the model, so one might think that these loops will dominate the simulated response to a disruption. But this is not the case. Delays and constraints limit the responsiveness of these loops in the days following a disruption.

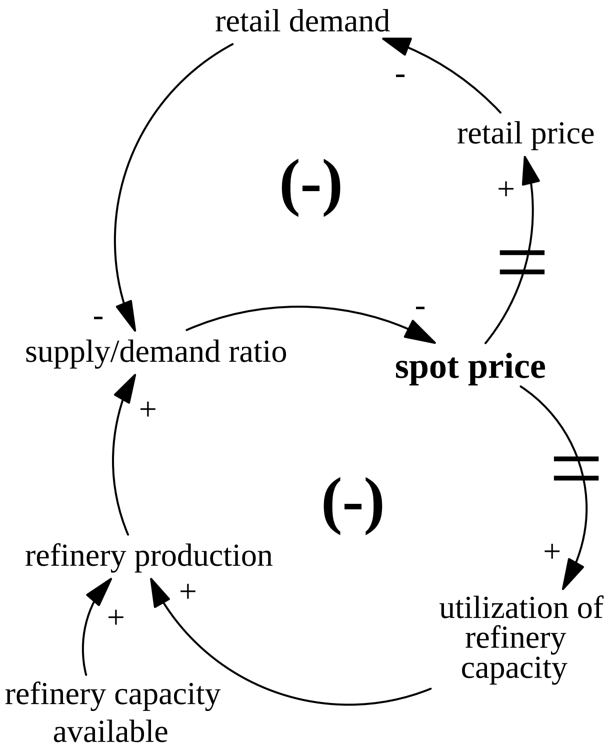

Figure 13. The sluggish loops in the model.

Imagine what would happen if there were an increase in the spot price due to a disruption. An increase in the spot price would lead to an increase in the retail price, a reduction in the retail demand, an increase in the supply/demand ratio and a subsequent reduction in the spot price. This is a negative loop whose actions are slowed by the long delay in passing spot price changes to the retail level. The model assumes an 18-day lag for spot price increases to be passed to retail; the lag for spot price reductions is even longer.

The refinery response loop is the lower loop in Figure 13. An increase in the spot price would lead to an increase in the utilization of the remaining capacity, an increase in refinery production, an increase in the supply/demand ratio and a subsequent reduction in the spot price.

The refinery loop is slowed by a lag as refinery owners observe and react to changes in the spot price. More importantly, this loop is limited by the starting assumption that 99% of refinery capacity is already in operation at the start of the simulation. The loops in Figure 13 provide only a small response in the days following a disruption.

Figure 14 turns our attention to the feedback loops that are more capable of responding to a simulated disruption. The upper loop is called the “storage control loop;” it represents the behavior of those who own the gasoline in private storage. Imagine how the owners would react in if there is a reduction in refinery production, an increase in the excess demand, withdrawals from storage, a reduction in the gasoline remaining in storage, an increase in the marginal convenience yield and an increase in the net convenience yield. Let’s continue around the “storage control loop” as if the spot price were fixed. An increase in the net convenience yield would lead to an increase in the percent net basis, a sign that gasoline in storage has become more valuable. Owners would then be reluctant to release stored gasoline into the spot market, and their reluctance would slow the reduction of gasoline held in private storage.

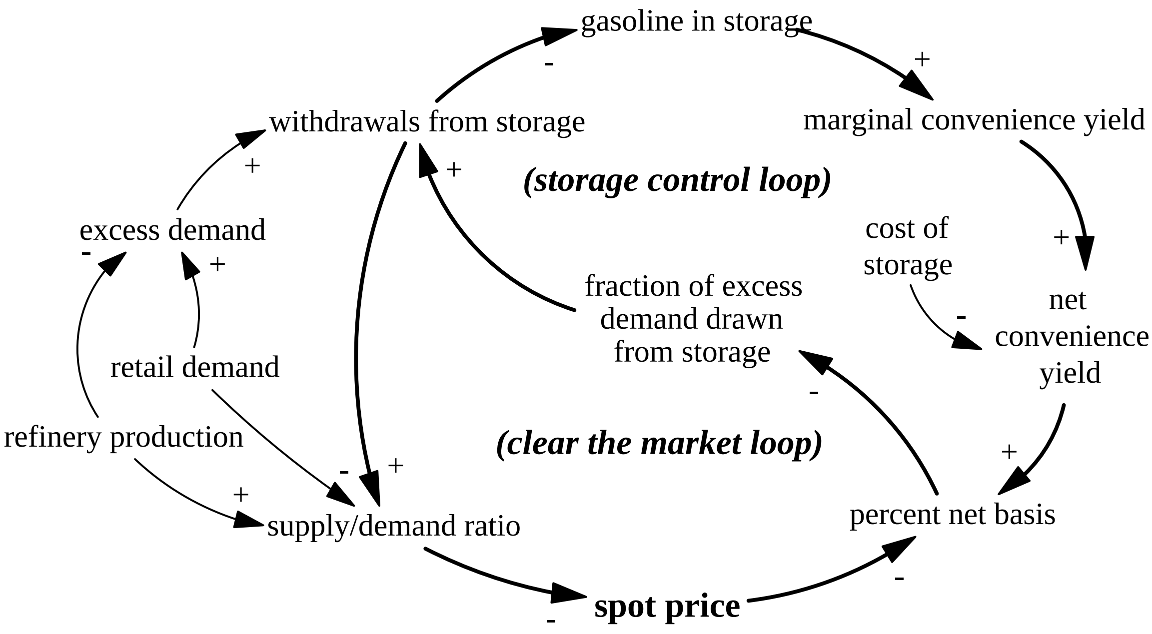

Figure 14. The feedback loops that are active in the days following a disruption.

But we know that the spot price does not remain fixed in the days following the disruption. It is driven upward by the actions of the "clear the market loop" in the lower part of Figure 14. This loop acts to raise the spot price whenever the TBD of supply falls below the TBD of demand. The market-clearing loop raises the spot price until sufficient supply is forthcoming to meet the demand. The main source of supply in the days immediately following the disruption is withdrawals from private storage. Consequently, the market-clearing mechanism acts to raise the spot price sufficient to induce owners of storage to release their gasoline into the spot market. This will happen when the spot price rises sufficiently to keep the percent net basis at or below the standard value.

The two loops in Figure 14 act continuously and without delay in the days following the disruption. The upper loop represents the behavior of owners of storage; the lower loop represents the market. Their combined actions will drive the spot price higher and higher until the disrupted refinery is back in operation or the emergency cargoes arrive from outside of California. With outages of short duration, the increases in spot prices will be capped when the refinery returns to operation. With outages of long duration, the increases in spot prices will be capped when the emergency cargoes arrive.

Impact of Different Disruptions

The illustrative simulation shown previously assumed a 15-day disruption of 100 TBD of refinery capacity. Figure 15 shows the impact of a 10-day disruption. It shows the spot price increasing during the 10 days that the refinery capacity is out of service. The spot price reaches a peak of 90 cpg in the 30th day and descends slowing over the next 6 days. The slow descent is made possible by the 100% utilization of the refinery capacity. With full utilization, the combination of refinery production and regular cargoes is somewhat larger than the retail demand. This causes a low slow decline in the spot price until the 36th day of the simulation. This day marks the arrival of emergency cargoes. Supply now exceeds demand by a larger margin, and the price descends more rapidly during this interval. The price falls below 60 cpg and gradually returns to 60 cpg during the second half of the simulation. The impact of the new disruption at $93 million, much lower than the $165 million impact reported for a 15-day disruption.

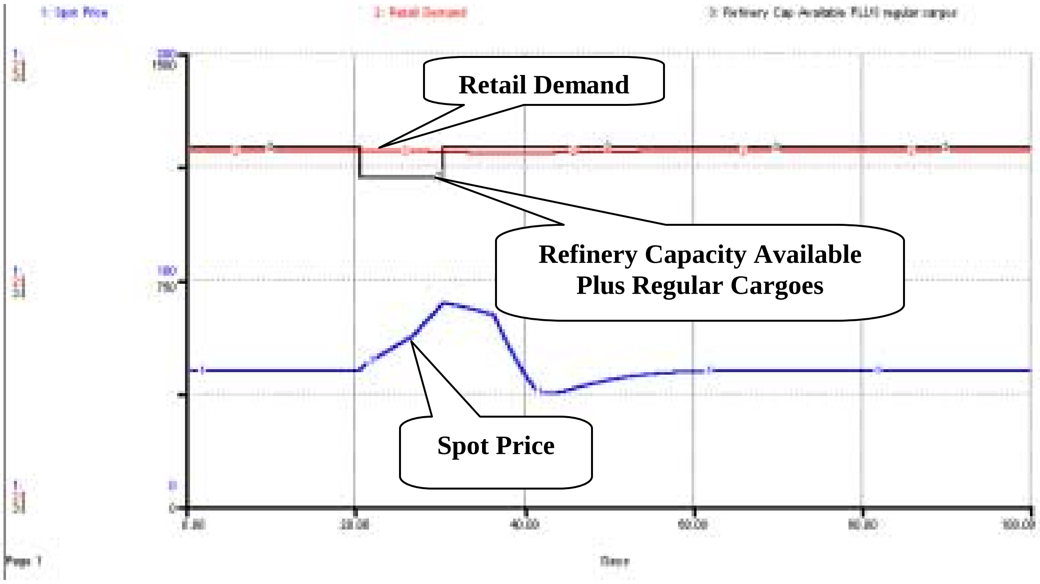

Figure 15. Simulated response to a 10-day disruption of 100 TBD of capacity.

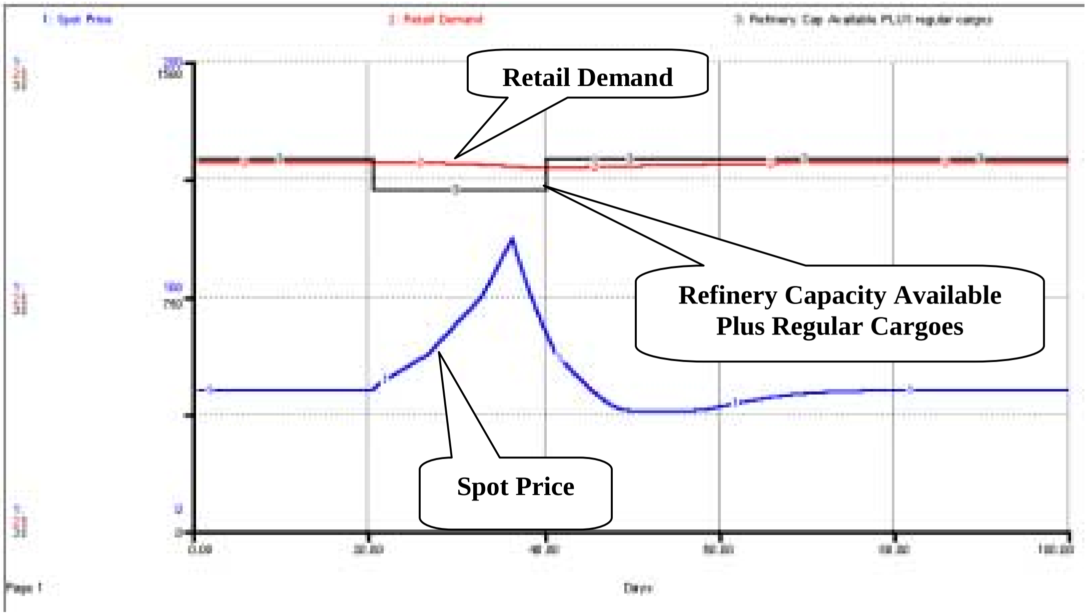

Figure 16. Simulated response to a 20-day disruption of 100 TBD of capacity.

Figure 16 shows the simulated response to a 20-day disruption of 100 TBD of refinery capacity. In this case, the refinery capacity does not return to operation until the 40th day of the simulation. The spot price would increase rapidly, peaking at 124 cpg by the 36th day of the simulation, the day when emergency cargoes begin to arrive in California. The spot price descends rapidly and bottoms out around 51 cpg by the 50th day of the simulation. The remainder of the simulation shows a slow recovery to equilibrium conditions. The simulated impact of this disruption is $184 million, only around 12% higher than the $165 million impact for a 15-day disruption..

Figure 17 compares the spot prices from the three simulations with a loss of 100 TBD of refinery capacity. We learn that the extension from 10 days to 15 days leads to a major increase in the run-up of prices. The peak price is 90 cpg with the 10-day disruption, 117 cpg with the 15-day disruption. The peak is only slightly higher with the 20-day disruption. The impact of the longer disruption is limited by the arrival of emergency cargoes. Recall that orders for these cargoes begin around 2 days after the start of the disruption and that the cargoes arrive after a 14-day transit time. The total lag time for emergency cargoes is around 16 days in these simulations. With these assumptions, disruptions much longer than 15 days are not likely to lead to significantly higher impacts.

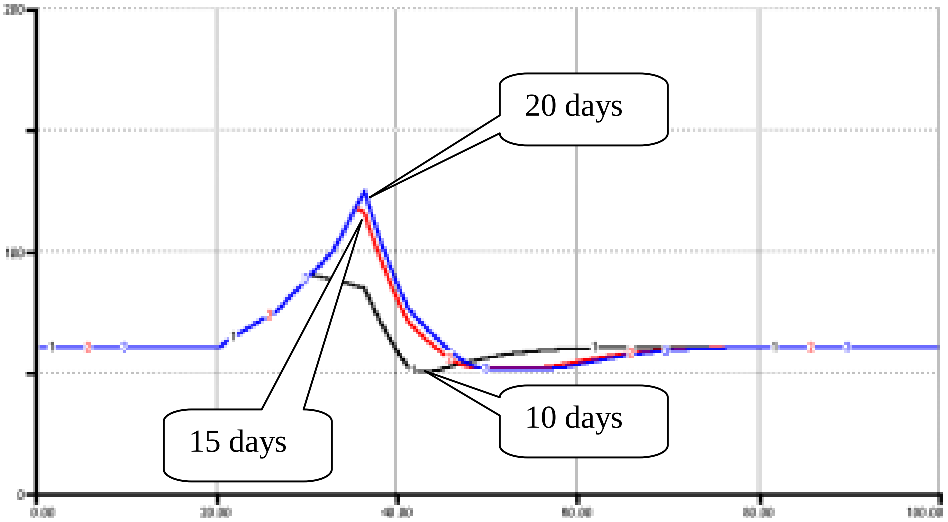

Figure 17. Spot prices with different durations of the 100 TBD disruption.

Figure 18 summarizes the impact in nine simulations with disruptions of different magnitudes and duration. The largest impacts appear in the back row for 150 TBD disruptions. These impacts would appear in the extraordinary case of a loss of 15% of the state's refinery capacity. This is a major loss, much larger than losses recorded by Finizza (2002, p 14). However, it is important to simulate these large disruptions because of their huge impact. (Indeed, the legislature emphasized the need to protect California from such disruptions when it called on the CEC to examine a disruption of California's largest refinery over a two-week period.) Finizza (2002, p 76) estimates the disrupted barrels for such an event at 2.3 million. Dividing 2.3 million barrels by 14 days gives the size of the disruption at 164 TBD. Thus, one may think of the "back row" results as indicative of the major disruptions which the legislature had in mind when calling for a study of the SFR. The “back row” impacts range from around $200 million to well over $500 million depending on the duration of the disruption.

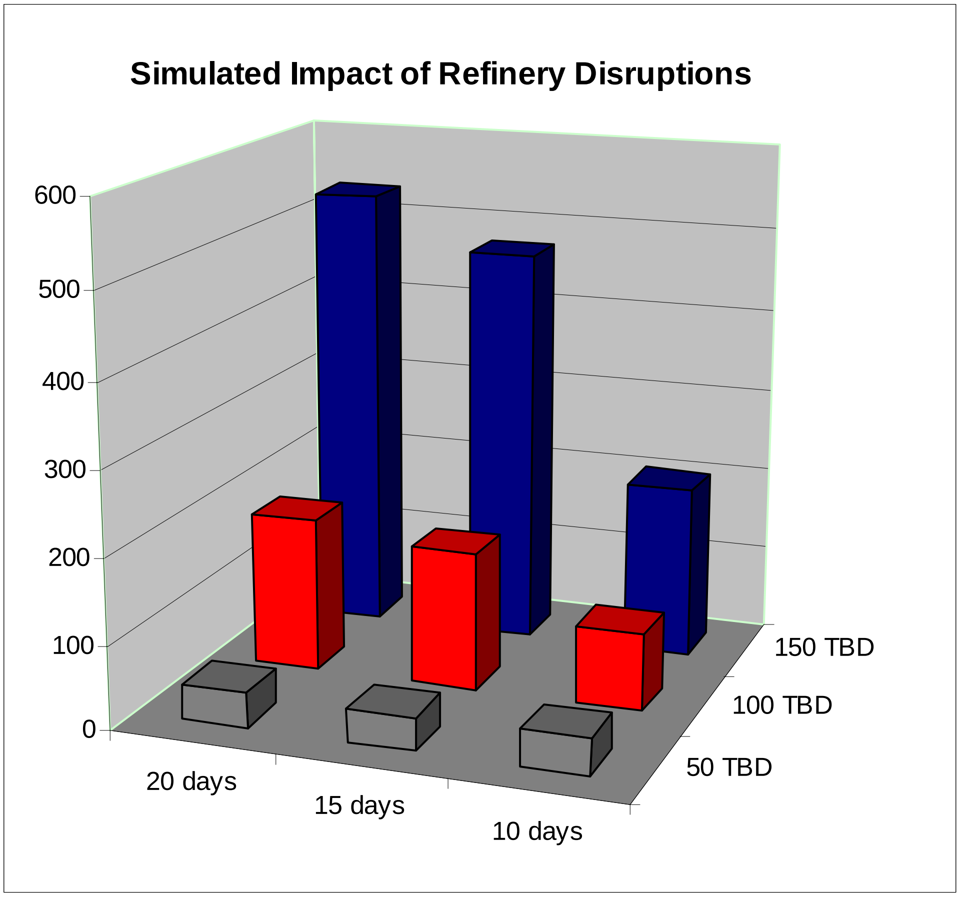

Figure 18. Simulated impact of nine refinery disruptions
(in millions of $ of increased wholesale payments for gasoline)

The middle row results in Figure 18 represent the loss of 10% of the state’s refinery capacity. This is also a major loss of capacity, and the model shows major impacts ranging from around $90 million to $180 million.

The front row results in Figure 18 represent the loss of 5% loss of the state’s refinery capacity. This is a smaller loss of capacity, and impacts are around $40 million. The results in the front row of Figure 18 may not seem as important, but smaller disruptions are much more frequent, according to the frequency distributions compiled by Finizza (2002, p 14). For example, the “average” disruption from Finizza (2002, p 13) is a 19-day outage of only 21 TBD of capacity. The disrupted barrels would amount to only about 0.4 million, somewhat smaller than the smallest disruption in Figure 18.

The Impact of a Strategic Fuels Reserve

Recall that Stillwater (2002, p 140) finds “overwhelming evidence” on the consumer benefits of an SFR and recommends the creation of a 5 million barrel reserve. Suggestions for the size of the reserve vary widely. Finizza (2002, p 77) estimates that a 0.9 million barrel reserve would cover all but 10% of refinery disruptions observed over 1996-2001. On the other hand, AB 2076 calls for an evaluation of a 2.3 million barrel reserve (Stillwater 2002, p 102). Figure 19 shows the variables that represent the operation of a SFR with the time swap mechanism described in the Stillwater report. The simulation begins with 5 million barrels of “Gasoline in the SFR” based on the Stillwater recommendation. The size is split between gasoline in the SFR and gasoline “on the water.” Figure 19 shows the stock of “Gasoline on the Water” to represent the gasoline that has been ordered by those with obligations to return gasoline to the SFR. Stillwater (2002, p 109) describes a situation where one might expect to see 1.5 million barrels of product on the water. (Companies with a return obligation can obtain the replacement gasoline anywhere they want, but the most likely source is cargoes shipped to California. So, in that sense, the gasoline is "on the water”.) Since we are using the model to simulate a single disruption, the simulations will begin with the stock of gasoline on the water set at zero; all 5 million barrels will be in the SFR.

The key flow in Figure 19 is the flow of “gas released from SFR.” This flow moves gasoline into a conveyor stock “Gasoline Transferred to Market” where the gasoline remains for a short interval before it contributes to “supply from SFR.” We assume that traders do not react instantaneously to changes in the spot price. Rather, it makes more sense to assume that traders watch the spot price over a “Price Observation Interval.” This lag time is set at one day, indicating that traders react to the smoothed, average value of spot prices over the previous day. We assume that traders compare the smoothed spot price to the “Average Value of the Traders’ Landed Cost for Return to SFR” to obtain the “Average Margin.” To illustrate, imagine that the spot price has climbed to 80 cpg in the days following a disruption, and the smoothed value of the spot price over the past day has reached 78 cpg. If a trader believes replacement supplies could be obtained for 68 cpg, the trader would expect a margin of 10 cpg if prompt supplies were obtained from the SFR. The average value of the traders’ landed cost is equal to the “Notional cost of External Supplies” a user input controlled by a slider on the model interface.

The “Notional Cost” is a term used by Finizza (2002, p 70) to represent the average value of the traders’ landed cost for return to the SFR. Page 71 of the Finizza report describes the notional cost of bringing Gulf Coast supplies to California at around 15 cpg above the Gulf Coast price. The cost premium is due to 10 cpg for transport and 5 cpg for meeting CARB requirements. In his example, the Gulf Coast cost was around 8 cpg below the typical spot price in California. This suggests that the notional cost is around 7 cpg higher than the regular spot price in California. If the starting value of the spot price is 60 cpg, the default value of the Notional Cost of External Supplies would be 67 cpg.

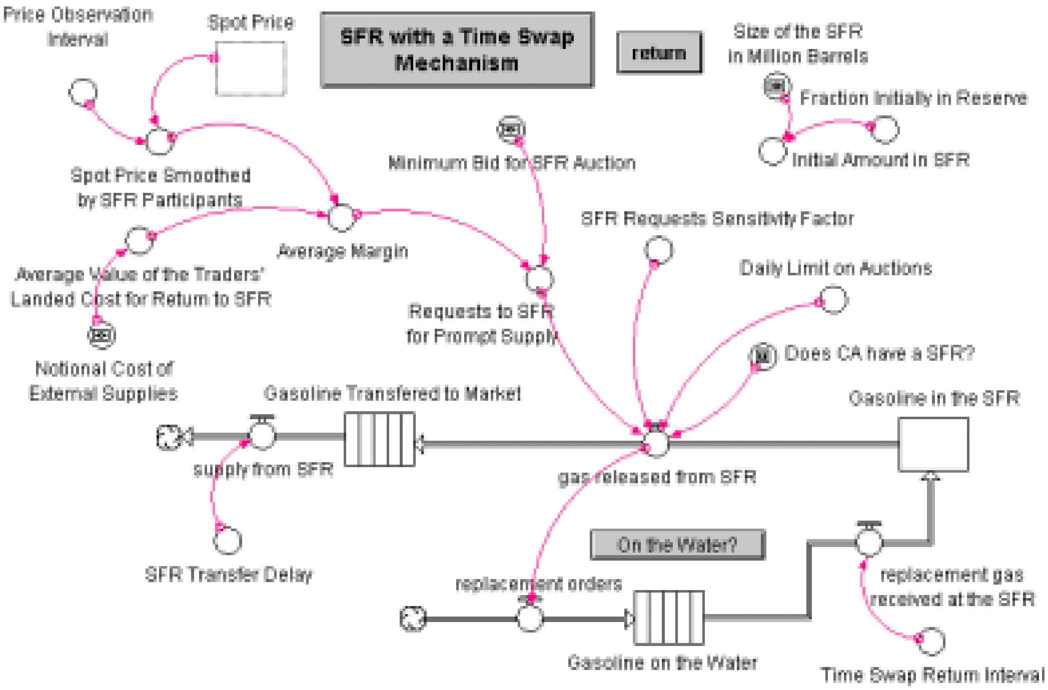

Figure 19. SFR Screen of the Model Diagram.

The final SFR parameter to be discussed in this paper is the "Fraction of Emergency Orders Displaced by a SFR." Recall that emergency orders are issued in the days following a disruption. For a 10-day disruption of 100 TBD of refinery capacity, for example, the emergency orders would amount to 1 million barrels of gasoline. Recall that these orders begin about 2 days after the outage and are spread over a 5-day interval. The gasoline arrives after a 14-day transit delay. However, the availability of prompt delivery from the SFR might cause some reduction in emergency orders following a disruption. A default value of 0.5 was selected as a plausible estimate of the "displacement fraction" based on discussions with CEC staff and consultants. In the example of a 10-day disruption of 100 TBD, for example, the emergency orders would only amount to 0.5 million barrels.

To illustrate the impact of a SFR, let's consider the 15-day disruption of 100 TBD of refinery capacity, the example described in detail previously. Figure 20 shows the spot price if this simulation is repeated with a SFR. The spot price would reach the notional cost within two or three days after the outage begins in the 20th day. After around 4 days, the spot price would exceed the notional cost by the 2 cpg required as the minimum bid for a SFR auction. Figure 20 shows that the spot price would continue to increase until around the 28th day of the simulation. This upward trend continues for several days because of the delays for traders to observe and smooth the spot price and the delay for gasoline released from the SFR to reach the market. When the gasoline does reach the market, it contributes to supply, and daily supply exceeds the daily demand. The peak price in this simulation is 79 cpg, far below the 117 cpg shown previously without a reserve.

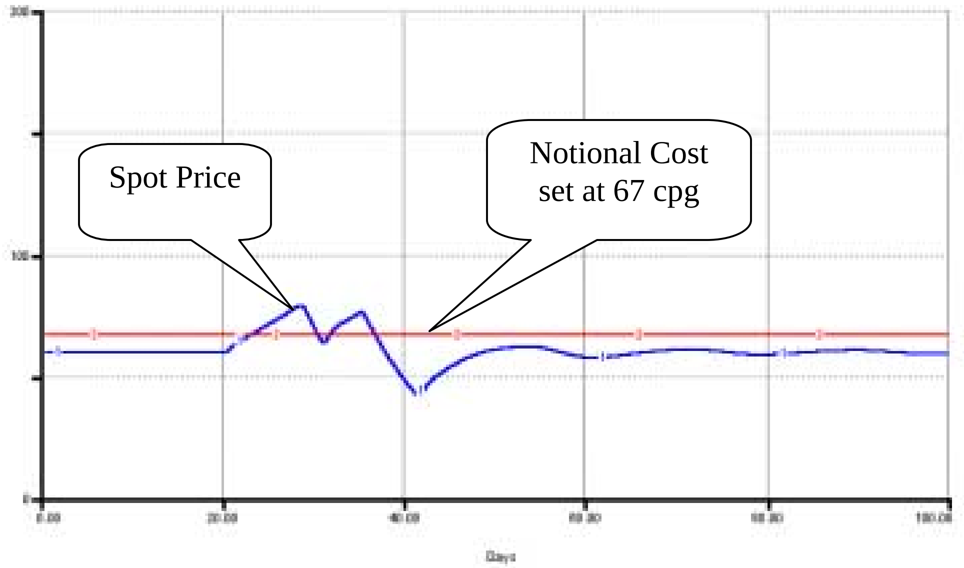

Figure 20. Simulated spot price with a 15-day disruption of 100 TBD with a SFR.

When the SFR supplies reach the market, they cause the daily supply to exceed the daily demand, and the spot price is driven downward. Within a few days, the spot price has fallen below the notional cost. But the outage continues for 15 days, and emergency cargoes have still not arrived in California. With these conditions, the spot price begins to increase again around the 31st day of the simulation. Within a few days, the spot price again exceeds the notional cost, triggering another round of auctions for prompt supply from the SFR. Figure 20 alerts us to the possibility that the presence of a SFR will not necessarily "cap" the spot price exactly at the notional cost. But the release of prompt supplies certainly reduces the overall increase in spot prices compared to the simulation without a SFR.

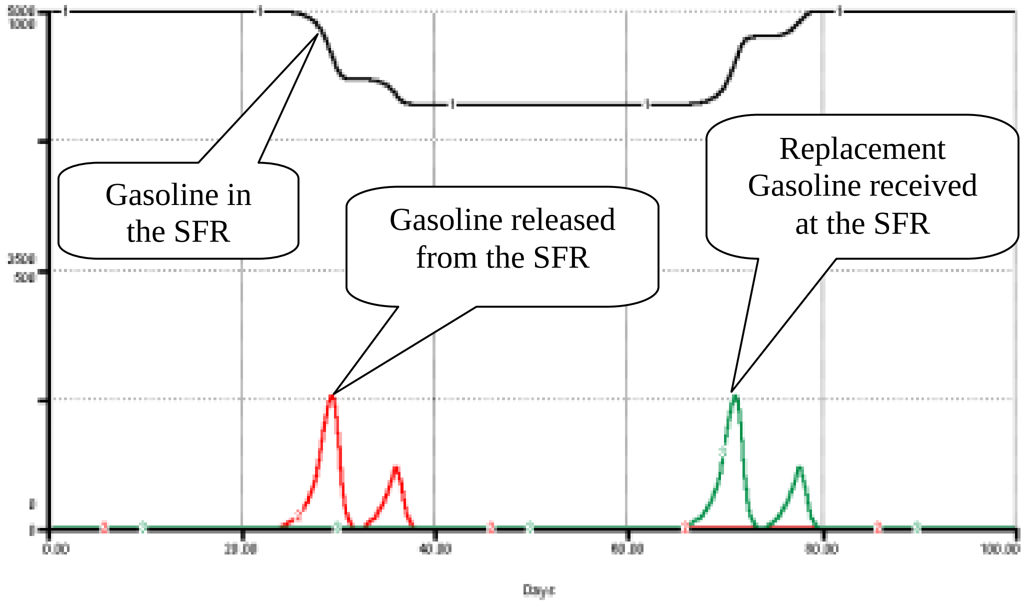

Figure 21. Gasoline released from the SFR with a 15-day disruption of 100 TBD.

Figure 21 helps us check whether there is enough gasoline in the SFR to deal with the disruption. The simulation begins with 5 million barrels of gasoline in the reserve. The reserve falls to around 4.3 million barrels after the initial round of auctions and to around 4.1 million barrels with the second round. The reserve returns to 5 million barrels when the replacement supplies are received after a 6-week delay.

Figure 22 compares the spot prices in the simulations with and without the SFR. The spot prices are identical during the first 25 days of the simulation. The comparison shows that the SFR leads to substantially lower prices from the 26th to the 45th days of the simulation. This benefit is achieved by the release of gasoline into the market. On the other hand, the SFR leads to somewhat higher spot prices after the 45th day of the simulation. Notice, for example, that the spot price has returned to around 60 cpg by the 50th day of the simulation. This is around the time that the spot price 'bottoms out" in the simulation without the SFR. This difference arises from the assumption that the presence of the SFR displaces around 50% of the orders for emergency cargoes.

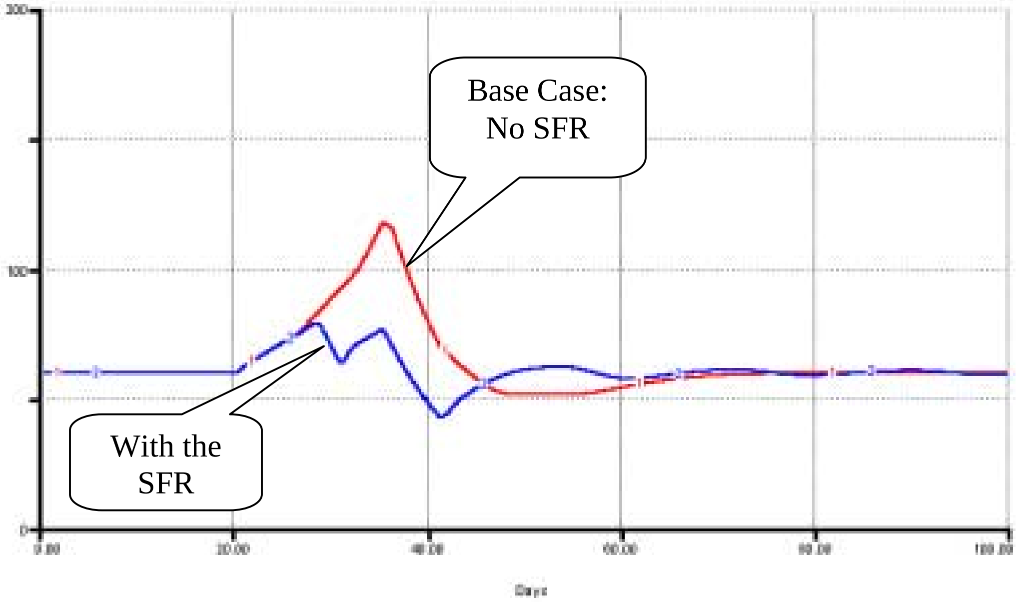

Figure 22. Spot prices in simulations of a 15-day disruption of 100 TBD of capacity.

Recall that our measure of the overall impact of a simulated disruption is the increase in cumulative wholesale payments for 40 million gallons/day of gasoline during the 100-day simulation. With the SFR, the impact is simulated $36 million, far below the $165 million impact in the simulation without the SFR.

Figure 23 shows impact from nine simulated disruptions when the model is operated with a SFR. (The $36 million impact described previously appears as the middle bar in the middle row of the bar chart.) This bar chart is arranged to allow for direct comparison with the bar chart for impacts without a SFR. Comparing Figures 18 and 23 shows that the most dramatic benefits of the SFR appear with the 150 TBD disruptions in the back row of the diagram. These benefits range from around $150 million to around $530 million depending on the duration of the disruption.

The middle row of Figure 23 shows impacts from a 100 TBD disruption with a SFR. They range from around $40 million to $60 million, depending on the size of the disruption. The front row of Figure 23 shows impacts from a 50 TBD disruption with a SFR. These impacts range from around $50 million to $60 million, depending on the size of the disruption.

The front row results for a disruption of 50 TBD of capacity alert us to the possibility that a SFR could lead to negative impacts. For example, the impact of a 20-day disruption is simulated at $52 million with the SFR, $42 million without the SFR. The presence of the SFR increases the impact of the disruption by $10 million. The negative benefit arises from the assumption that the presence of the SFR will cause some displacement of the emergency cargoes ordered in the days following the disruption.

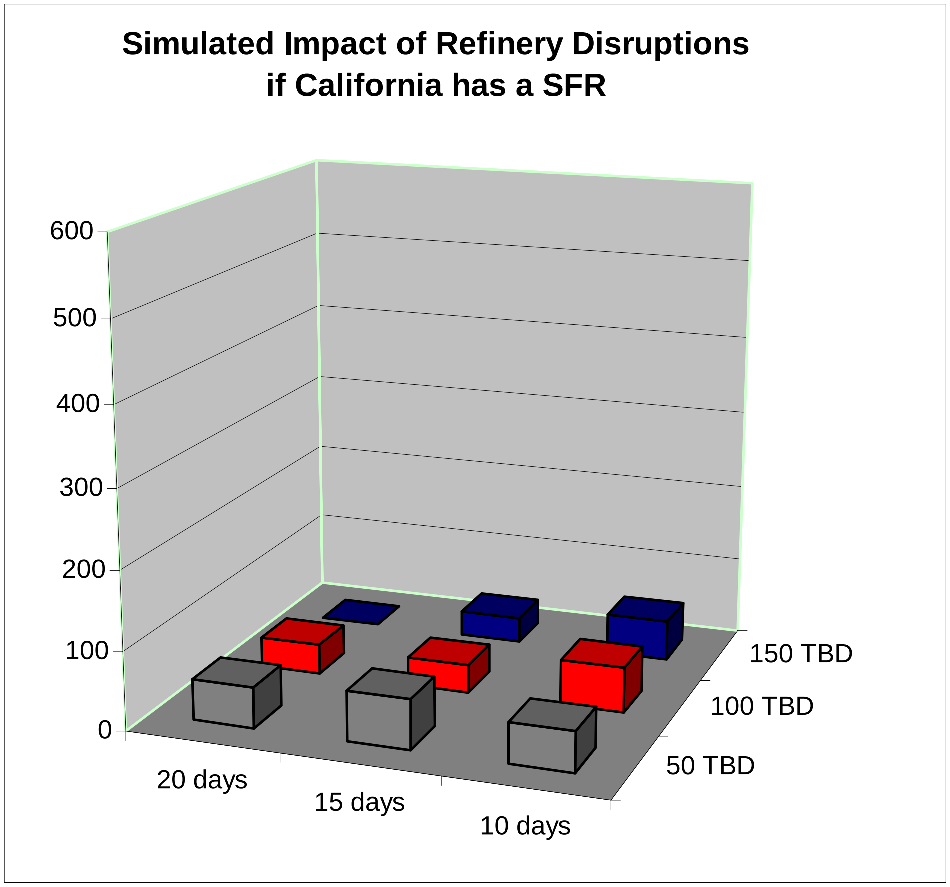

Figure 23. Simulated impact of nine refinery disruptions if California has a SFR (in millions of $ of increased wholesale payments for gasoline.)

The reduction in emergency orders is one example of displacement that analysts keep in mind when anticipating the impact of the SFR. Displacement might also appear in the form of a reduction in private storage. According to the Attorney General report (AG 2000, p 10), for example, the subcommittee assigned to study reserves “expressed concerns that a state reserve would lead refiners to reduce their own inventory levels.” On the other hand, the Stillwater report (2002, p 70) states that it is “very unlikely that the presence of a reserve limited to only 2.3 million barrels and designed with a release mechanism that creates forward liquidity, will have any significant impact on inventories currently held by the industry." The Finizza report (2002, p, 80) recommends an "intense analysis of private versus public storage to ascertain the possibility of crowding out" (of private storage by public storage).

Future Work

It is important to emphasize that the analysis in this paper is illustrative, not definitive. The purpose has been to illustrate the use of system dynamics to simulate the impact of a SFR under a variety of different disruptions. A definitive analysis of the SFR would be conducted by the CEC staff. The CEC is looking at a variety of changes to develop the model. Model development is an ongoing process of construction, testing, discussion and further development. The goal is to develop a model with greater realism, recognizing along the way that realism is a relative concept. That is, we should compare the realism of the current model with the realism of whatever alternative model might be used to estimate the impact of a disruption.

In my opinion, the most important task for further model development is to experiment with the model to compare simulated spot prices with the spot price increases following an actual disruption. CEC staff should also compare the simulated draw down of private storage with the draw down following an actual disruption. This is an extremely challenging task as the data on disruptions is complicated by a wide variety of factors which make it difficult to select a "typical disruption" suitable for comparison. When the CEC builds its confidence in simulating the benefits of an SFR for individual disruptions, it should repeat the type of simulations shown here for a wide variety of disruptions of different magnitude and duration. The simulated benefits could then be combined with estimates of the frequency of the different disruptions to obtain an expected, overall benefit of the SFR.

From the illustrative results in this paper, one would expect the SFR to deliver huge benefits for large disruptions. For example, the benefit of a SFR in mitigating the price impact of a single disruption (15-day disruption of 150 TBD) is estimated at over $400 million. This estimate for a single event is approximately the same as the $400 million in estimated benefits for an entire year that appear in the Finizza and Stillwater reports. With this huge benefit, one would anticipate the expected value calculation to be dominated by the frequency of large disruptions. Interestingly, these are the type of disruptions envisioned by AB 2076 when it called on the CEC to consider a disruption of the largest California refinery over a two-week period.

Finally, it is important to emphasize that the SFR is but one of several options under consideration in California. Other options include the creation of a public information system (on inventories and shipments) or the provision of incentives to increase the amount of gasoline held in private storage. The system dynamics model described in this paper allows such options to be examined through interactive simulation.

References

AG 2000	Attorney General Bill Lockyer, "Report on Gasoline Pricing in California," May 2000.
Duffie 1989	Darrell Duffie, Futures Markets, Prentice Hall, 1989.
Finizza 2002	Anthony Finizza, "Economic Benefits of Mitigating Refinery Disruptions: A Suggested Framework and Analysis of a Strategic Fuels Reserve," CEC Report P600-02-018D, July 2002.
Ford 1999	Andrew Ford, Modeling the Environment: An Introduction to System Dynamics Models of Environmental Systems, Island Press, 1999.
Keynes 1930	John Maynard Keynes, Treatise on Money, Volume II, (The Applied Theory of Money), Macmillan and Company, 1930.
Kindleberger 1978	Charles Kindleberger, Manias, Panics and Crashes: A History of Financial Crises, Basic Books, New York, 1978.
Pindyck 2001	Robert S. Pindyck, "The Dynamics of Commodity Spot and Futures Markets: A Primer," The Energy Journal, Vol 22, No. 3. 2001.
Stillwater 2002	Greg Haggquist, David Hackett and Thomas Gieskes, "California Strategic Fuels Reserve," study conducted by Stillwater Associates, CEC report P600-02-017D, July 2002.
Sterman 2000	John Sterman, Business Dynamics, Irwin McGraw-Hill, New York City, 2000.
Working 1949	H. Working, "The Theory of Price of Storage," American Economic Review, Vol 39, 1949, pp 1254-1262

Back to the Top