International Petroleum Production Model (IPPM)

March, 2008

Table of Contents

Overview/Background 1
Model Inputs 2
  Global World Oil Balance Excel File 2
  Miscellaneous Excel File 2
Core Model Concepts 3
  Supply Curves 3
  Production Profiles 5
Model Calculations for the Historical Period 6
  Resolution of Historical Data 6
  Beginning Recovery Factor and Supply Curve Calibration 7
Flowchart of Model Forecast Calculations 7
Model Forecast Calculations for the Current Year 8
  Calculate a New Recovery Factor Limit Based on Price Changes 8
  Shift Supply Curves for Technology Change 9
  Calculate Production from New and Existing Capacity 9
  Calculate Finding and Development Cost for New Capacity 9
  Calculate Net Present Value for Revenue Stream for New Capacity 10
  Calculate Profit for New Capacity 11
Model Forecast Calculations to Build New Capacity for the Next Year 11
  Determine Needed Overall Production for Next Year 11
  Build New Capacity in Small Increments to Maximize Profit 11
  Production Quotas and Production Growth Limits 12
  Tie Some New Capacity Choice to Recent History 13
  No New Capacity Allowed Beyond the Upper Limit 14
  No New Capacity Allowed if the Profit from it is Negative 14

Overview/Background

This documentation describes the calculations made in the Fortran coded version of the International Petroleum Production Model (IPPM). The model determines international petroleum production starting in 2008 and continuing to 2200, from resources in five regions and differentiated by six production streams. The model relies upon a variety of input data and assumptions for each stream and region that are described in more detail later. Some of the key inputs include information about initial resource in place and the ultimate recovery factor along with the production profiles and the resource supply curves.

The model "builds" new capacity in each region and for each stream by maximizing profits based on net present value of the finding and development costs (from the resource supply curves) and the revenue streams. In each year the decision as to where to add the new capacity to meet the needed production is broken down by the model into a large number of small increments. In each increment the model solves for the current stream and region with the maximum profit and then moves that stream/region up its supply curve. The overall solution is the sum over all the small incremental decisions.

The IPPM currently includes five regions (although it can easily be expanded to more) consisting of:

1 - OPEC – Saudi Arabia
2 - OPEC – Other Middle East
2 - Other OPEC
3 - United States
4 - Other Non-OPEC

The IPPM currently includes six production streams consisting of:

1 - Crude and Condensate
2 - Natural Gas Plant Liquids (NGPL)
3 - Extra Heavy Crude
4 - Bitumen
5 - Shale
6 - Source Rock

The model does all its calculations in billions of barrels per year. Some of the inputs are in millions of barrels per day and are converted to billions of barrels per year. The initial year for the model is 2008 and the forecast horizon extends to 2200.

Model Inputs

Global World Oil Balance Excel File

This file is a large repository of the Global World Oil Balance (GWOB) data and assumptions. Only a small portion of the information in this file is used. This includes:

Initial In Place – The original amount of resource that exists in each region and for each stream in billions of barrels.
Recovery Factor Limit – The fraction of the initial in place resources in each region and for each stream that is ultimately recoverable.
Production – Historical data in each region and for each stream from 1980 to 2007 in millions of barrels per day.
Cumulative Production – Historical data in each region and for each stream in 2007 in millions of barrels per day. This is the sum of all past production.
Real Price – Forecast of world petroleum price from 2008 to 2200 in dollars per barrel. This is a scenario assumption.
World Petroleum Consumption – Forecast of total world petroleum consumption from 2008 to 2200 in millions of barrels per day. This is a scenario assumption.
Technology Change – Forecast of indices of technology change for each stream from 2008 to 2200. This is a scenario assumption.

Miscellaneous Excel File

This is a smaller Excel file that contains additional data and assumptions. All of the information in the file is used. This includes:

Supply Curve Parameters – These include an initial cost for the beginning recovery factor, a final cost associated with the recovery factor limit, and a coefficient that determines the shape of the curve. These are designated for each region and for each stream. The supply curves are explained in more detail below.
Production Profiles – These designate the vintages of production for new capacity (fraction of the new capacity that can be produced over its lifetime). These are designated for each region and for each stream. These are explained in more detail below.
Production Quotas – These designate the maximum amount of production in any one year for each region and for each stream in millions of barrels per day.
Production Growth Limit – These designate a limit on the amount of growth in production over the previous year. These are designated for each region and for each stream and can vary for each forecast year.
Production Growth Limit Lower Bound – An lower bound on the amount of production that is allowed before there is a growth limit. These are designated for each region and for each stream.
Increments - Number of decision increments to use for making new capacity choices.
Decision Switch - Switch for choice of whether to use least cost or maximum profit for making new capacity choices
Technology Switch - Switch for whether or not to use technology change to shift supply curves.
Technology Change Parameters – These include an overall coefficient and a coefficient applied to the low price for each region and for each stream.
Lease and Operating Cost – Designated for each region and for each stream in dollars per barrel and used in the net present value calculation.
Taxes – Designated for each region and for each stream as a fraction and used in the net present value calculation.
Cost of Capital – Designated for each region and for each stream in dollars per barrel and used in the net present value calculation.
New Choice Historical Fraction – The fraction of the new capacity choice decision that uses the new capacity choices in the previous year (or years) versus being based upon the least cost or maximum profit economic choice. This can vary by forecast year.
New Choice Number of Previous Years – The number of previous years used as the history when a fraction of the new choice is based upon history.

Core Model Concepts

Supply Curves

The IPPM uses supply curves for each of the five regions and for each of the six streams to represent the relationship between the finding and development (F&D) cost for new reserve capacity and the cumulative level of proved developed producing (PDP) reserves. The cumulative level of PDP reserves is the sum of cumulative production in the previous year plus the current level of PDP reserves. On the supply curve, the cumulative level of PDP reserves on the X-axis is actually represented as the ratio (or fraction) of the cumulative PDP reserves relative to the initial in place (IIP). This fraction is the same as the recovery factor (RF). The supply curve is defined for costs associated with recovery factors that run between the current beginning level and the ultimate recovery factor. The ultimate recovery factor is the fraction of initial in place (IIP) that can be ultimately recovered.

Basic Supply Curve. In the first year of the model horizon, the beginning recovery factory (RF), or fraction of the current beginning cumulative PDP reserves relative to IIP for any one stream and region might be something like 0.15. The F&D cost associated with this beginning RF might be something like $10 per barrel. These two values would then be the beginning lower left-hand value on the supply curve for that stream and region.

An exogenous assumption used by the model is that there is an ultimate recovery factor for each region and for each stream. In other words, in most cases not all of the IIP can be ultimately recovered; there is an upper limit which is a smaller fraction of the IIP. This recovery factor limit for any one stream and region might be a fraction such as 0.60 indicating that only 60% of the IIP can be ultimately recovered. An additional assumption might be that the F&D cost associated with this recovery factor limit might be something like $250 per barrel. These two values would then be the ending upper right-hand value on the supply curve for that stream and region.

The supply curve now consists of beginning lower left-hand values and ending upper right-hand values. These points need to be connected in some fashion. One option would be to fill in a large number of intervening points and interpolate between them. Another option is to draw a line between them using a power function with a coefficient that will provide a reasonable shape. We have used the second option for its computational ease and for its functional characteristics. The functional form of the supply curve is given by:

Cost = LowCost + Alpha * (RF* ^ Beta) ,

where: Alpha = (HighCost - LowCost) / (RFRange ^ Beta) .

First, in these equations we have assumed that the RF* that is used has been shifted to go from 0.0 to the total of its entire range. In other words, the starting RF above of 0.15 is shifted to 0.0 and the RF limit of 0.60 is shifted to 0.45, which then becomes the RFRange. If we assume that the value of Beta is 2.0 in our example, then Alpha is equal to 1185.2. Using Alpha, we can calculate the F&D cost at a recovery factor of 0.20 (shifted to 0.05) which equals $12.96 per barrel.

Technology Shifted Supply Curve. Due to technology changes the supply curve will shift downward, changing the relationship between the F&D cost and the current recovery factor. For the same recovery factor, the cost of developing new capacity will be less. In the model, technology change is represented by an index that decreases as technology improvements increase. The index is then used to shift the beginning lower left-hand cost and the ending upper right-hand cost for the supply curve and the curve is redrawn between these two points. There are two coefficients that can be used in association with the technology indicies. The first coefficient is the fraction of the overall technology index that is used and the second coefficient is the additional fraction of the index that should be used for the lower cost point. Using the example above, assume that the technology index is 0.9 and the first coefficient is 1.0 and the second coefficient is also 1.0. The lower cost then becomes $9 per barrel and the upper cost becomes $225 per barrel. A new Alpha is calculated and is 1066.7. Solving for the same recovery factor of 0.20 as above (shifted to 0.05) the F&D cost now becomes $11.67 (instead of $12.96).

Examples of the Basic Supply Curve and the Technology Shifted Supply Curve as described in the examples above are shown in the following chart. The way that the supply curves work is also documented in an Excel spreadsheet named "SupplyCurveDemo.xls" that is separately available.

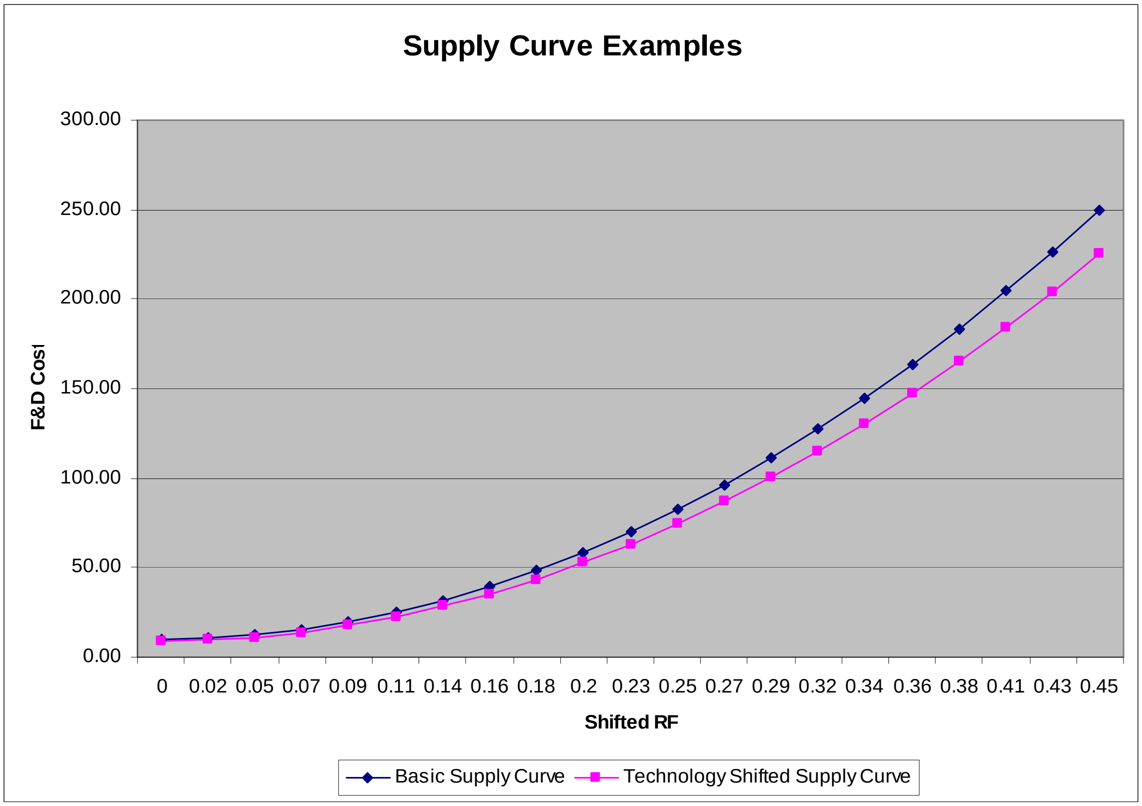

Production Profiles

In the IPPM it is assumed that new capacity will be produced over a period of time according to a specific production schedule that is given by a production profile. The production profiles used by the model can vary for each of the five regions and for each of the six streams. Generally speaking, the production profiles in the IPPM are for a period of 25 years or vintages (they can be shorter, but cannot be longer), so that all of the new proved developed producing (PDP) capacity built for the current year has been produced by the end of 25 years. Since the production profile shows how the capacity is used and exhausted over the period of 25 years, the sum of all the fractions making up the production profile must add up to 1.0.

The following table shows an example of a production profile for one region and one stream in which 8% of the new capacity is produced in the first year. So, for example, if there is 10 billion barrels of new capacity added that is first available in 2010, then in 2010 there will be 0.8 billion barrels produced from this new capacity. In 2011, there will be 0.7 billion barrels produced. This capacity will continue producing over the 25 years so that in 2034 there will be 0.3 billion barrels produced from that new capacity that became available in 2010. In 2034 that will be the last of the production from this capacity and in 2035 there will be no more production for that new capacity that was added in 2010. The model keeps track of each piece of new capacity added in each year for each region and for each stream, and produces from it over the lifetime given by the production profile.

Vintage	Production		Vintage	Production		Vintage	Production

	Profile			Profile			Profile
1	0.08		10	0.04		19	0.03
2	0.07		11	0.04		20	0.03
3	0.06		12	0.04		21	0.03
4	0.05		13	0.04		22	0.03
5	0.04		14	0.04		23	0.03
6	0.04		15	0.04		24	0.03
7	0.04		16	0.03		25	0.03
8	0.04		17	0.03			
9	0.04		18	0.03			

A key concept is that in any one year, the total production for each region and each stream is the sum of production from each piece of new PDP capacity that was previously added in the 25 previous years. An example for one region and stream is shown in the following table. The table shows how the total production in a year such as 2030 consists of production from the previously added new capacity. The total production in 2030 is the sum of the production in each of the 25 vintages, shown in the last column (for brevity, several rows of vintages are not shown).

Total Production in 2030 is the Sum of Production from Previously Added New Capacity				
Year Added	**Vintage**	**New Capacity**	**Production Profile**	**Production by Vintage**
2030	1	10.0	0.08	0.80
2029	2	9.5	0.07	0.67
2028	3	10.3	0.06	0.62
2027	4	9.4	0.05	0.47
2026	5	8.9	0.04	0.36
2025	6	8.7	0.04	0.35
…	…	…	…	…
2007	24	3.9	0.03	0.12
2006	25	4.2	0.03	0.13

Another key concept concerns how new capacity is added. If new capacity is being added in order to satisfy a production need in the year it becomes available, much more new capacity needs to be added than the amount of needed production. For example, consider that capacity is being added in 2020 in order to satisfy a production need of 2 billion barrels per year. Only a small fraction of the new capacity can be used for production in the first year, according to the production profile. Using the production profile above, only 8% of the new capacity can be used to satisfy production in the first year. Therefore, if we need production of 2 billion barrels per year, we must add 25 billion barrels of new capacity (calculated as 2 / 0.08 = 25).

Model Calculations for the Historical Period

Resolution of Historical Data

The first forecast year for the IPPM is 2008. In that year the model starts producing based upon the capacity additions for proved developed producing (PDP) that have been added in the past (over the 25 years of the production profiles). The historical data provides information on the amount of production in previous years, but does not provide information on the amount of PDP over the previous 25 years. However, using the production data over the period for 1980 through 2007 along with the cumulative production in 2007 and the assumed production profiles, we can construct the rest of the data. The data we calculate includes cumulative production in each year, the PDP new capacity in each year, and the cumulative PDP capacity in each year.

Cumulative Production in Each Historical Year. In each region and for each stream we start with cumulative production in 2007 and the stream of production for each year going back to 1980. Cumulative production in 2006 is simply cumulative production in 2007 minus the production in 2007. This calculation can be continued back to 1980.

Proved Developed Producing (PDP) New Capacity in Each Historical Year. In each region and for each stream we start with the production in 2007 and the production profile. We know that the production in 2007 was produced from new PDP capacity previously added for the current year and over the previous 24 years and we know how this production is scheduled over the vintages based upon the production profile. But there are still too many unknowns. We have to make a starting assumption about how much new capacity was added in each of those years. We know that it probably changes over time, but as our first guess we assume a level in one year and that it grows at the same rate as production has grown. Once we have our first guess at all the values we can then solve over time and determine the resulting production for 2007. We can then compare this result to the actual production in 2007 and use that factor to adjust the new PDP capacity in all the years by the same factor. This, then, is our estimate of the stream of new PDP capacity over the previous 25 years. This allows us to run the model off of a internally consistent set of data. (Note that these do not go all the way back to 1980, they go back 25 years.)

Cumulative PDP Reserves in Each Historical Year. In each region and for each stream, the cumulative PDP reserves is the sum of the cumulative production from last year plus the PDP reserves in the current year. As we did for cumulative production we can work backwards to get cumulative PDP reserves over the historical period.

Beginning Recovery Factor and Supply Curve Calibration

In this model, the recovery factor (RF) is the ratio of the cumulative PDP reserves to the initial in place (IIP). In the first year of 2008, each region and each stream have a level of cumulative PDP reserves and an IIP, so the beginning RF can be calculated for each.

As described above, the supply curve is defined as running between the beginning RF and the RF defined as the upper limit on the fraction of IIP that can be produced. The beginning RF is calculated above and the RF limit is based on user assumptions. The costs associated with each of these along with the supply curve coefficient are inputs to the model. At this point it is necessary to calculate an “Alpha” calibration factor for each of the supply curves in each region and stream. This was detailed in the above discussion of the supply curves.

Flowchart of Model Forecast Calculations

The model begins execution in 2008 and continues to loop through the years until 2200. It performs its calculations in each year taking into account the five regions and the six streams.

Some of the calculations are made for the current year and some are made in the current year but looking ahead to the next year. The following flowchart shows the major activities, and they are described more fully in the following sections.

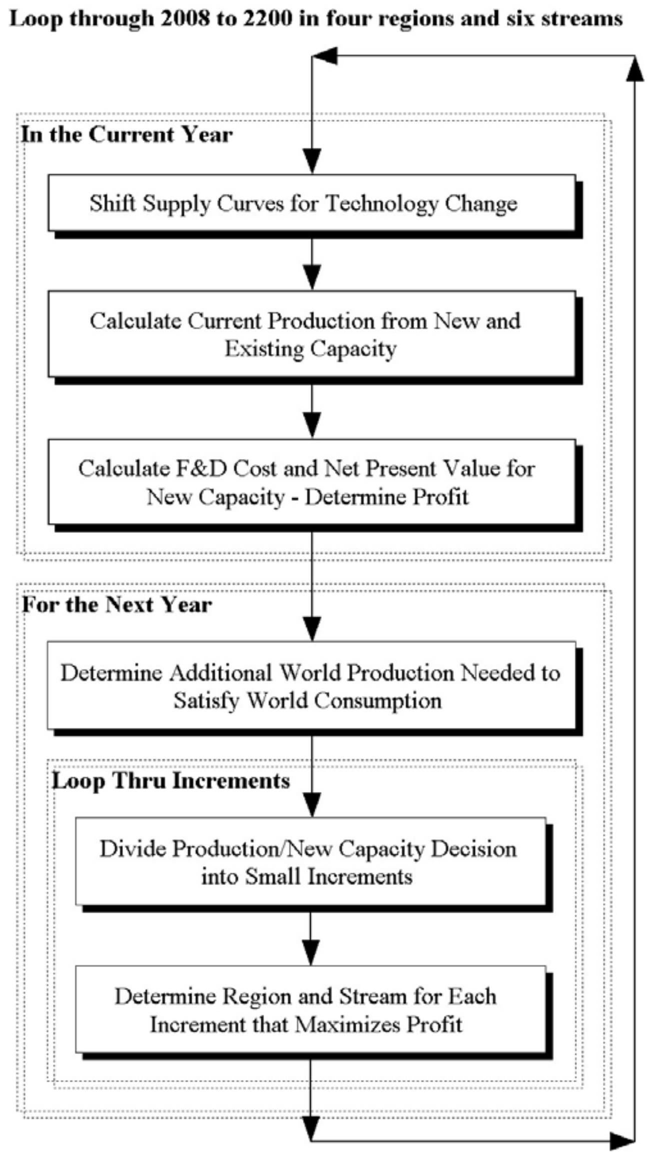

Model Forecast Calculations for the Current Year

There are several things that the model calculates in the current year. These include shifts to the supply curves for technology change, the amount of production from new and existing capacity, and the finding and development cost and the net present value for new capacity.

Calculate a New Recovery Factor Limit Based on Price Changes

In each forecast year, the model recalculates the recovery factor (RF) limit based upon how prices have changed from the previous year. The amount of the change is typically small and is based upon user specified coefficients. The idea is that as the world oil price increases, the RF limit will be somewhat higher. As explained earlier, the RF limit is used in conjunction with the initial in place (IIP) to determine the ultimate amount of resource that is recoverable (cumulative production). It is also used to define the upper limit on the supply curve which is recalculated each year (see below).

The recalculation of the RF limit changes the value from the previous year relative to how prices have changed from the previous year. The RF limit and the coefficients are defined for each of the six streams and five regions.

RFLim(y) = RFLim(y-1) + (1.0 + RFLim(y-1)) *
((RFLim(y-1) * (Price(y) / Price(y-1)) ^ RFCoef) – RFLim(y-1))

where: RFLim is the recovery factor limit,
Price is an exogenous petroleum price in dollars per barrel,
RFCoef is a user selected coefficient.

Shift Supply Curves for Technology Change

In each forecast year, technology change will shift the supply curve down, allowing the same amount of new capacity at a lower cost in each region and for each stream. The first thing that the model does in each forecast year is to shift the supply curve based upon the change in the technology index. Note that this shift also takes into account the change in the RF limit described above. How the technology change shift works was detailed in the above discussion of the supply curves. Once the supply curve has been shifted, it is necessary to re-calibrate it by calculating a new value for Alpha.

The technology change index is defined for each of the six streams. The value for this index is exogenous to this model and is basically a user assumption.

Calculate Production from New and Existing Capacity

In each forecast year, the total amount of production is calculated in each region and for each stream. This is simply based upon the amount of new capacity that was added in this year and previous years and on the production profile. This process was described above in the discussion of production profiles and is the sum of all the vintages of production:

TotalProduction(y) = Sum over 25 vintages [NewCap(v) * ProductionProfile(v)]

The total production is the sum of the production from each of the pieces of new capacity added in this and in previous years. For this reason, even if no new capacity was added this year, there can still be production from new capacity added in previous years.

Calculate Finding and Development Cost for New Capacity

The finding and development (F&D) cost for the new capacity that was put into place in the current year for each region and stream is based upon the supply curve. The amount of new PDP capacity is added to the previous year's cumulative production to get the current cumulative PDP reserve level. This is divided by the initial in place (IIP) to get the current recovery factor (RF).

The supply curve is solved for the current recovery factor to find the F&D cost. This cost is the marginal cost because it is the cost on the margin, or the cost at the right-hand side for the total amount of the cumulative PDP reserve.

F&DCost = LowCost + Alpha * (CurRF ^ Beta) ,

where: F&DCost is the finding and development cost in dollars per barrel,
LowCost is the lower cost on the supply curve,
Alpha is the calibration coefficient for the supply curve,
CurRF is the current recovery factor,
Beta is the coefficient for the supply curve.

Calculate Net Present Value for Revenue Stream for New Capacity

The calculation of the net present value (NPV) of all the future discounted revenue streams for the new capacity added involves a number of different factors. The basic idea is that this new capacity will provide a level of production over a number of years into the future. The amount of production over the future years is based upon the production profile. In the examples above, this new capacity will produce 8% of its total in the current year, 7% in the next year, 6% in the next year, on through the 25 years of the production profile. In each of these years there is also a petroleum price, lease and operating costs, taxes, and tax deductions for finding and development costs. There is also a cost of capital that provides a rate at which the future revenues are discounted. The revenue in each of the future years is given by:

Revenue(fy) = [(Price(fy) – LOC + TaxDed(fy)) * (1 – TaxRate)] * Production(fy)

where: Price is an exogenous petroleum price in dollars per barrel,
LOC is the lease and operating cost in dollars per barrel,
TaxDed is the tax deduction in dollars per barrel (calculated below),
TaxRate is a tax rate in fractional terms,
Production is the production in billions of barrels.

The production value in the equation above is the amount of production from the new capacity in the current year over each of its future years of production based upon the production profile:

Production(fy) = NewCap(y) * ProductionProfile(fy)

The tax deduction in the equation above is the deduction for the finding & development cost depreciated over the lifetime of the production. It is depreciated at the same rate as production, in other words it is based on the production profile as:

TaxDed(fy) = F&DCost(y) * ProductionProfile(fy)

The revenue above is the revenue in each of the future years. It is necessary to discount this revenue in future years based upon how far out each year is and the cost of capital:

DiscountFactor(fy) = 1 / [(1 + COC) ^ (fy – y)]

The total net present value of the overall revenue from the new capacity is the sum over all the years of the revenue in each year times the discount factor in each year:

TotalNPV(y) = Sum over all years [Revenue(fy) * DiscountFactor(fy)]

This number represents the current value of all revenue for all of the production from the new capacity. We are also interested in determining the net present value per barrel of capacity so that it can be compared directly to the F&D cost which is in dollars per barrel. This is simply the total current value of revenue divided by the barrels of new capacity:

NPV(y) = TotalNPV(y) / NewCap(y).

Calculate Profit for New Capacity

The profit for the new capacity in the current year is simply the net present value minus the F&D costs:

Profit(y) = NPV(y) – F&DCost(y) .

Model Forecast Calculations to Build New Capacity for the Next Year

In the current year the model also does planning for next year by determining how much new capacity to build in order to satisfy future production needs and by determining how to build that new capacity in order to maximize profits.

Determine Needed Overall Production for Next Year

In the current year, the model needs to determine how much new capacity it needs to build for next year in order to satisfy production needs in the next year. The model assumes that it will provide just enough production in order to meet the total world consumption needs. Total world consumption is exogenous to the model and is based upon the user scenario. In the current year the model is able to determine how much potential production can be produced next year based upon the existing, previously-built capacity. This is calculated the same way as total production in the current year, except that there is no current year and the production profile is shifted by one year.

PotentialProduction(y+1) = Sum over 24 vintages [NewCap(v) * ProductionProfile(v)]

Once the model has determined the potential production for next year for each of the steams and regions it adds it up to get the total world potential production for next year. It then compares this total world potential production next year to the world consumption next year to determine how much additional production will be needed next year.

NeededProduction(y+1) = WorldConsumption(y+1) – PotentialProduction(y+1)

We do not yet know how much new capacity needs to be built because we don't yet know the makeup of the new capacity in terms of which regions it is built in and which streams are used. The model will make those decisions in order to maximize profit.

Build New Capacity in Small Increments to Maximize Profit

At this point we know how much new production will be needed next year. However, we do not yet know the sources of that new production. Each potential source (from any of the five regions or using any of the six streams) for that new production has potentially a different finding and development cost and potentially a different net present value for returns from the investment.

The model could solve for the F&D cost from the supply curves and could solve for the NPV in each of the five regions and for each of the six streams. It could then determine the profit in each of those 24 categories and find the category with the maximum profit. It could then determine the total amount of new capacity that would have to be added in that region for that stream in order to provide the needed production. However, this is a very lumpy and sub-optimal way of solving for where to build the new capacity in order to meet the needed production.

Economic Decision Increments. Instead, the model breaks down the economic decision process for determining where to add new capacity into a number of small decision increments. The number of decision increments is user selected and typically the model has been run using 25 increments but it has often been run using 100 increments. Basically, what the model does is take the total needed production and breaks it down into much smaller production increments. Assuming we are using 25 increments, then:

IncProduction(y+1) = NeededProduction(y+1) / 25 .

The model then goes through the 25 increments of production and determines the region and stream for which each subsequent piece of the incremental production can be added with the maximum profit. Each time it solves for the region and stream with the maximum profit it keeps track and then moves on to the next increment of production. However, since the previously chosen region and stream has now added more capacity and has moved up its supply curve, it can now add the next increment only at a higher cost. For this reason, in this next increment, a different region and stream may have the maximum profit and be chosen. The greater the number of increments the less lumpy the results and the finer the resolution of chosen categories. We are typically using 25 increments because that seems to provide us the resolution we need.

As the model goes through each of its decision increments, a F&D cost is calculated for each region and stream based upon its supply curve. However, the supply curve is based upon the concept of cumulative production and a recovery factor as described above. So it is necessary to determine how much new capacity would have to be built in each of the regions and streams in order to provide the incremental new production. This new capacity will likely vary among the regions and streams because the production profiles vary among regions and streams. As shown earlier, the new capacity needed to provide a given production for the first year depends upon the first vintage of the production profile:

IncNewCap(y+1) = IncProduction(y+1) / ProductionProfile(v=1)

So if we are adding incremental production of 0.1 billion barrels and one category has a production profile with a first vintage value of 0.8 and another has a value of 0.10, in the first instance we need 1.25 billion barrels of new capacity and in the second instance we need 1.00 barrels of new capacity. So in the first case we go up the supply curve by a greater amount than in the second case. It is important for the model to keep track of these different changes so that it can make sure that the movements along the supply curves are accounted for properly as we go through the decision increments.

Production Quotas and Production Growth Limits

As the model builds new capacity to meet needed production, there might be some constraints on how much production might be allowed from some regions and streams. The model has two types of these constraints, one a limit on the amount of production from a region and stream, and the second a limit on how fast production from a region and stream can grow. Both of these are user determined and it is a user option as to whether to use one or the other or neither.

An example of a production quota is the limit on production set by OPEC. There is an input file to the model in which the user can set these quotas for any years and for any region and stream. Then in each year the model proceeds to add production in an incremental fashion and choose the region and stream with the maximum profit. However, in each year the model keeps track of the cumulative production for each region and stream as it goes through the decision increments. If there is a production quota for some region and stream and it goes over the quota at some point during the incremental process, then that region and steam is not allowed to add any new production in subsequent increments. In subsequent increments, even if this region and stream has lower costs the model will only choose from the other regions and streams. Since the production is not stopped until after the region and stream has gone over its quota, it is possible for the quota to be exceeded by a small amount (this depends upon the size of the increments).

Production growth limits would typically be used for regions and streams that might suddenly become cost effective and add large amounts of production beyond infrastructure capabilities. There is an input file to the model in which the user can set the growth limits for any years and for any region and stream. This works in much the same way as the quotas, except that the model limits the production to a fraction of the production in the previous year. In this way, the model limits the rate of growth, but after a number of years the absolute amount could have expanded considerably.

A problem that can occur when a production growth limit is put on some stream and region that previously had zero production, is that it will effectively never allow there to be any production in that stream and region. In this model we get around that problem by placing a lower bound on the production level that is subject to the growth limit. In other words, if the lower bound is set at 0.25 billion barrels, then the stream and region can add up to this amount before it becomes subject to the growth limit. This lower bound can be input for each region and stream so that the model will not use the growth constraint as long as the production is below this level.

Tie Some New Capacity Choice to Recent History

As described to this point, the model bases its decision on regions and streams for the new capacity choice in each year on purely economic decisions. The model looks at the finding and development costs from the supply curve and the net present value of the future revenue streams and makes the economic decision based upon the maximum amount of profit. There may be a number of reasons why the pure economic model may not be the best choice. These may include that there is some habit to doing what was done in the past, the infrastructure has already been put into place, there is a value to prior knowledge and prior experience, and there might be costs that are less and are not being captured in the supply curves (for example, developmental wells versus exploratory wells).

The IPPM allows the user to make a portion of the new capacity choice decision based upon what was done in the previous year (or years). The remaining portion of the decision would then be based on the pure economics. In other words, if the history decision fraction were 0.25, then the model would look at the previous year’s new capacity choices and how those are allocated among the regions and streams and use that same allocation for the first 25 percent of the new capacity added in this year. The remaining 75 percent of the new capacity added would then be allocated based upon the incremental economic decision process that was described earlier.

One exception is that the model still looks at the quotas and growth limits. If any of the regions and streams exceeds its quota or growth limit based upon the historical allocation, then it is set only to its limit. The amount over its limit is added back into the portion that is allocated based upon the incremental economic decision process.

Another feature is that the model allows the user to use more than one previous year of new capacity additions to determine the allocation. Using any one year could tend to be somewhat volatile, so the model allows the user to use up to ten previous years. In other words, if the user selected three previous years then the model would choose the historical allocation based upon the allocation of new capacity choices in the previous three years (instead of just one year).

No New Capacity Allowed Beyond the Upper Limit

The model will not allow any new capacity to be added in any decision increment that puts that stream and region above the upper limit. The upper limit in this case is new capacity that adds production so that the overall cumulative production from that stream and region exceeds the recovery factor (RF) limit times the initial in place (IIP). This is a fairly basic concept that is at the core of the model. The assumptions behind the model are that there are upper limits to production for any stream and region, but the model has to explicitly enforce that limit. This limit applies to all new capacity, both that added based upon history and that added based upon the economic decision increments.

No New Capacity Allowed if the Profit from it is Negative

The model will not allow any new capacity to be added in any decision increment if that new capacity brings a negative profit. In the choices based upon the economic decision increments the model chooses the stream and region for each increment based upon maximum profit. Any stream and region in which the profit is negative is not allowed to be chosen. This also applies to the part of the choice decision that is based on history. An economic choice decision is not involved when a portion of new capacity is added based upon the choices in prior years. However, for that portion of new capacity, the model will calculate costs, revenue streams, and the profits for each stream and region and will not allow that new capacity to be built if the profit is negative. That capacity is instead put into the pool of new capacity that will be chosen using the incremental economic decisions.

This means that if there is no region or stream that has profit, then none of them will be chosen and the model will not be able to add enough new capacity to meet the production needed in the next year. In other words, there will not be enough economically viable production in order to satisfy the exogenous level of demand at that price level.