Closed Loop Control with Jump Processes: OPEC Oil Production Cases Study

Jad Soucar
Department of Industrial and Systems Engineering
University of Southern California
Los Angeles, CA 90089, USA

1 Proposed Model

1.1 Setting & Assumptions

OPEC is comprised of 13 oil producing nations that coordinate their efforts to manipulate global oil prices through increasing or decreasing production. The group must constantly contend with changing market conditions and global demand conditions to adjust production accordingly. Sometimes member nations fail to meet their production quotas, while other times member nations intentionally produce and export more oil due to localized economic and financial pressure. While the interplay between member countries along with its intrinsic game-like dynamics are an interesting area of research in and of itself, we focus our attention on an idealized case. First we assume that OPEC nations are able to coordinate and adjust their strategy instantaneously to respond to market conditions. Second we assume that all member countries are able to meet their production quotas or production cuts. Third we assume that if an OPEC nation fails to meet their obligations other member countries will not explicitly retaliate and will instead produce the remainder of the oil necessary to reach the quota. Finally we assume that any oil produced will be immediately purchased and consumed so long as their is demand for that oil.

1.2 Dynamics

Under these assumption we can define the rules of our game. First the dynamics of the global price of oil at time P_t depend on global demand D_t along with the total barrels of oil being produced $q_t \geq 0$.

$$dP_t = P_t[\eta \log\left(\frac{D_t + \epsilon}{q_t + \epsilon}\right)dt + \sigma_P dW_t^P] \tag{1}$$

Where η represents the sensitivity of price to demand and production, The log term produces negative drift when the production surpasses demand and there is a surplus, and creates a positive drift when there is a shortage. ϵ is set to some small value to avoid numerical instability. We note that $\sigma_P W_t^P$ is a Brownian motion corresponding to any small scale fluctuations in price that are caused by market pressures external to our system.

$$dD_t = D_{t^-}[(\alpha_D - \beta P_t - \lambda k)dt + \sigma_D dW_t^D + (e^Y - 1)dN_t] \tag{2}$$

Next we analyze demand D_t dynamics in equation 2. First we claim that, by standard microeconomic theory, D_t will shrink when P_t rises and grow when P_t declines by some rate β. We also assume a base level of demand growth α_D, which encodes industrialization trends and population growth. We also claim that demand is often hit by non-periodic shocks of various sizes. For example the recession of 2008 and the Coronavirus pandemic both resulted in over 70% decreases in demand. To that end we incorporate dN_t into our dynamics, which represents the number of shocks within time dt drawn from a Poisson distribution with rate λ. $Y \sim \mathcal{N}(\mu, \delta^2)$ is the log-size of each jump. Note that we use the demand immediately before t (i.e D_{t^-}) since a jump at t would result in a discontinuous process. We let $\sigma_D dW_t^D$ be small scale fluctuations in demand caused by market pressures external to our system. Finally we define $k = \mathbb{E}[e^Y - 1]$ so we don't double count shocks in both the drift and dN_t term.

$$dX_t = (\mu - \lambda k)dt + \sigma dW_t + (e^Y - 1)dN_t \tag{3}$$

We note that dD_t is almost identical to the Merton Jump Diffusion Model where the state is modeled as equation 3[1]. Now that the relevant state dynamics have been defined we can define the probability space that we are working in. Namely $(\Omega, \mathbb{F}, (\mathcal{F}_t)_{t=0}^T, \mathbb{P})$, where $\mathcal{F}_t$ is right continuous filtration generated by (W_t^D, W_t^P, N_t, Y_i) and $\mathbb{P}$ is the joint distribution of all the Brownian, Normal, and Poisson distributions. We now look towards the value function of the oil producer.

$$V(D, P, t) = \sup_q \mathbb{E}^{D,P}\left[\int_0^T e^{\rho(t-T)}(\min\{q_t, D_t\}P_t - C(q_t))\right] \tag{4}$$

Next we define the value function $V(D, P)$ which seeks to maximize the discounted expected profit, where ρ is the discount rate. Note that $\min\{q_t, D_t\}P_t$ is the revenue at time t assuming that if there is demand for oil it will be sold immediately after being produced. We also introduce a cost of production function $C(q_t)$ to de-incentivize endless production. Specifically we let $C(q_t)$ be the following where c_1 is the baseline cost of production per barrel of oil while c_2 is the marginal penalty for scaling production beyond current operational capacity. Note that we assume that any supply produced beyond what is demanded is immediately discarded. This treats oil as perishable which we concede isn't realistic but simplifies our model. Finally we note that the terminal reward is 0, meaning that all profits must be accumulated during the the time horizon $[0, T]$.

$$C(q_t) = c_1 q_t + c_2 q_t^2 \tag{5}$$

Next we derive the associated Hamiltonian-Jacobi-Bellman Partial Integro-Differential Equation (HJB PIDE) associated with the previously defined closed loop control problem. To that end we first find dV_t using the Ito formula for jump processes [2]. In our case $dV(D_t, P_t, t)$ is the following, where $V^- = V(D_{t^-,P_t,t})$ and $d\tilde{D}_t = D_{t^-}[(\alpha_D - \beta P_t - \lambda k)dt + \sigma_D dW_t^D$ is the continuous terms in dD_t.

$$\begin{aligned} dV_t =& \partial_t V^- dt + \partial_D V^- dD_t + \partial_P V^- dP_t \\ &+ \frac{1}{2}[\partial_{DD}^2 V^- (d\tilde{D}_t)^2 + \partial_{PP}^2 V^- (dP_t)^2] \\ &+ \partial_{DP}^2 V^- d\tilde{D}_t dP_t \\ &+ [V(D_{t^-}e^Y, P_t, t) - V(D_{t^-}, P_t, t) - (e^Y - 1)D_{t^-}\partial_D V^-]dN_t \end{aligned} \tag{6}$$

We note that this formula for dV is fairly intuitive since it captures the change in V with respect to each input (t, D_t, P_t), the change due to Brownian shocks, along with the change due to jumps beyond first order effects. The last term is functionally the residual of a first order Taylor expansion, Since the ∂_D term already captures first order effects of the jump process. This mirrors the form of the pure diffusion Ito formula. Next we use the infinitesimal generator dV_t to derive the associated HJB PIDE by leveraging the dynamic programming principle.

$$\begin{aligned} V(D,P,t) &= \sup_{q_t} \mathbb{E}^{\mathbb{P}}\Big[\int_t^{t+h} e^{-\rho(s-T)}(\min(q_s, D_s)P_s - C(q_s))ds + e^{-\rho h}V(D_{t+h}, P_{t+h}, t+h)\Big] \\ \xrightarrow{h\to 0} 0 &= \sup_{q_t} \mathbb{E}^{\mathbb{P}}\Big[e^{-\rho(t-T)}(\min(q_t, D_t)P_t - C(q_t))dt + dV_t\Big] \end{aligned} \tag{7}$$

We expand equation 7 by taking the expectation $\mathbb{E}^{\mathbb{P}}[dV_t]$. Note that $\mathbb{E}^{\mathbb{P}}[dW_t^{D/P}] = 0, \quad \mathbb{E}^{\mathbb{P}}[dN_t] = \lambda dt, \quad (dW_t)^2 \sim dt$, and recall that $k = \mathbb{E}^{\mathbb{P}}[e^Y - 1]$.

$$\begin{aligned} \mathbb{E}^{\mathbb{P}}[\partial_t V^- dt] &= \partial_t V^- dt \\ \mathbb{E}^{\mathbb{P}}[\partial_D V^- dD_t] &= \mathbb{E}[\partial_D V^- d\tilde{D}_t] + \mathbb{E}[\partial_D V^- D_{t^-}(e^Y - 1)dN_t] \\ &= \partial_D V^- D_{t^-}(\alpha_D - \beta P_t - \lambda k)dt + \partial_D V^- D_{t^-}\lambda \mathbb{E}[e^Y - 1]dt \\ &= \partial_D V^- D_{t^-}(\alpha_D - \beta P_t)dt \\ \mathbb{E}^{\mathbb{P}}[\partial_P V^- dP_t] &= \partial_P V^- P_t \eta \log\left(\frac{D_t + \epsilon}{q_t + \epsilon}\right)dt \\ \mathbb{E}^{\mathbb{P}}\left[\frac{1}{2}\partial_{DD}^2 V^- (d\tilde{D}_t)^2\right] &= \frac{1}{2}\partial_{DD}^2 V^- D_{t^-}^2 \sigma_D^2 dt, \\ \mathbb{E}^{\mathbb{P}}\left[\frac{1}{2}\partial_{PP}^2 V^- (dP_t)^2\right] &= \frac{1}{2}\partial_{PP}^2 V^- P_t^2 \sigma_P^2 dt, \\ \mathbb{E}^{\mathbb{P}}\left[\partial_{DP}^2 V^- d\tilde{D}_t dP_t\right] &= \partial_{DP}^2 V^- D_{t^-} P_t \sigma_D \sigma_P \rho dt \\ \mathbb{E}^{\mathbb{P}}[(\dots)dN_t] &= \lambda\left(\mathbb{E}_Y\left[V(D_{t^-}e^Y, P_t, t)\right] - V(D_{t^-}, P_t, t) - \partial_D V^- D_{t^-}k\right)dt. \end{aligned}$$

Putting this all together we get an exact form of $\mathbb{E}^{\mathbb{P}}[dV_t]$. Recall that $Y \sim \mathcal{N}(\mu, \delta^2)$ is the expected jump magnitude so $k = \mathbb{E}_Y[e^Y - 1] = e^{\mu + \frac{1}{2}\delta^2} - 1$.

$$\begin{aligned}\mathbb{E}[dV_t] = \Bigg[&\partial_t V^- + \partial_D V^- D_{t^-}(\alpha_D - \beta P_t) + \partial_P V^- P_t \eta \log\left(\frac{D_t + \epsilon}{q_t + \epsilon}\right) \\ &+ \frac{1}{2}\partial^2_{DD} V^- D^2_{t^-}\sigma^2_D + \frac{1}{2}\partial^2_{PP} V^- P^2_t \sigma^2_P + \partial^2_{DP} V^- D_{t^-} P_t \sigma_D \sigma_P \rho \\ &+ \lambda\left(\mathbb{E}_Y\left[V(D_{t^-}e^Y, P_t, t)\right] - V(D_{t^-}, P_t, t) - \partial_D V^- D_{t^-} k\right)\Bigg] dt\end{aligned} \tag{8}$$

We can use equation 8 to rewrite the associated HJB PIDE of equation 7 as the following. Note that we set the terminal reward to 0, however the terminal condition can just as easily be some function $\psi(D_T, P_T, T)$ which penalizes unmet demand, or any number penalization or reward schemes. Also note that the HJB PIDE is deterministic, so we can drop the dependence on D_{t^-} and write $\mathbb{E}_Y$ in integral form where ϕ is the pdf of Y.

$$\begin{aligned}0 =& \partial_t V + \partial_D V D_t(\alpha_D - \beta P_t) + \\ &+ \frac{1}{2}\partial^2_{DD} V D^2_t \sigma^2_D + \frac{1}{2}\partial^2_{PP} V P^2_t \sigma^2_P + \partial^2_{DP} V D_t P_t \sigma_D \sigma_P \rho \\ &+ \lambda\left(\int_{-\infty}^{\infty} V(D_t e^y, P_t, t)\phi(y)dy - V(D_t, P_t, t) - \partial_D V D_t k\right) \\ &\sup_{q_t}\left\{e^{-\rho(t-T)}(\min(q_t, D_t)P_t - C(q)) + \partial_P V P_t \eta \log\left(\frac{D_t + \epsilon}{q_t + \epsilon}\right)\right\} \\ V(D_T, &P_T, T) = 0\end{aligned} \tag{9}$$

2 Numerical Simulations

Now that we've formally derived the associated HJB PIDE we can apply numerical simulation schemes to solve for the the optimal revenue V and production strategy q within the time horizon $[0, T]$. The model we introduced above has quite a few parameters, so we first find appropriate parameters in literature. Then we choose to implement the Monte Carlo schemes to solve the HJB PIDE in equation 9. We also include the algorithm necessary to solve the problem using the finite difference scheme.

2.1 Setting Parameters

Recall that our model has 10 parameters, $\{\eta, \sigma_P, \alpha_D, \beta, \lambda, \sigma_D, \mu, \delta, c_1, c_2\}$. We will analyze each parameter and choose a reasonable value from literature.

We first focus on the parameters contained in the price dynamics 13. We first note that the supply $q_t \geq 0$ for all t. Next we have η which represents the sensitivity of oil price to shortages and surpluses. In macro-economic literature this parameter is referred to as the short-run elasticity of oil demand and is often times estimated to be anywhere between 5% to 10% annually. In other words a 1% imbalance in demand vs. supply can induce a 5% to 10% price change [3]. To that end we set $\eta = 1.075$ Next we look at σ_P which represents the volatility of oil prices. Empirically implied oil volatility is often between $0.25 \pm .05$ annually. [3].

Second we look at the parameters contained within the demand dynamics. First we have α_D which represents the baseline demand growth from populations growth and industrialization. Empirical estimates indicate that oil consumption grows about 1.3% to 1.7% annually [4]. Next is β, which represents the price effect on demand. Qualitative analysis indicates that only for large fluctuations in price is demand dampened. So we estimate a small $\beta \approx 0.01$ [3]. Finally we look at σ_D which represents the volatility in demand. Aside from discrete jumps in demand the volatility is estimated to be about 1% annually [3].

The jump process has three parameters that we must set. First we set $\lambda = .1$ annually, which represents a major demand disruption every 10 years [4]. the mean log jump size μ is set to -0.05 and the standard deviation $\delta = .1$ which would correspond to an average $e^Y - 1 \approx -5\%$ plunge in demand. Additionally these parameters are inline with recent demand shocks which almost always skew negative. With the largest negative and positive shocks in recent history being -9% in 2020 and 3.3% in 2010 [4].

The production cost function $C(q)$ is quadratic with two parameters c_1 and c_2. c_1 represents the baseline production cost per barrel which is estimated to be about \$10 per barrel, which the marginal cost c_2 is estimated to be about \$0.1 [5]. We summarize these parameters in the table bellow. We can treat ρ as a measure of inflation since cash flow today will be worth more then cash flow in the future. So to that end we set $\rho = 2\%$ which is the federal reserve's annual inflation target.

2.2 Monte Carlo Method

We first choose to solve the HJB PIDE for the optimal production strategy q_t using a simple Monte Carlo neural network approach. We let q_t be a trainable neural network $\pi_\theta(D_t, P_t, t)$ with parameters θ. We can discretize the dynamics using Euler-Maruyama with time steps Δt:

$$P_{t+\Delta t} = P_t \left(1 + \eta \log\Big(\frac{D_i + \epsilon}{q_{i,j}^{n*} + \epsilon}\Big)\Delta t + \sigma_P \sqrt{\Delta t}\xi_P\right),$$

$$D_{t+\Delta t} = D_t \left(1 + (\alpha_D - \beta P_t - \lambda k)\Delta t + \sigma_D \sqrt{\Delta t}\xi_D + (e^Y - 1)dN_t\right),$$

where $\xi_P, \xi_D \sim \mathcal{N}(0,1)$ are independent standard normals, $Y \sim \mathcal{N}(\mu, \delta^2)$, and

Parameter	Description	Annual	Monthly
η	Sensitivity of price to shortages/surpluses	1.075	1.075/12
σ_P	Price volatility	0.25	$0.25/\sqrt{12}$
α_D	Baseline demand growth	0.015	0.015/12
β	Demand sensitivity to price	0.001	0.001
λ	Jump intensity (shock frequency)	0.1	0.1/12
σ_D	Demand volatility	0.01	$0.01/\sqrt{12}$
μ	Mean log jump size	-0.05	-0.05
δ	Std dev of log jump size	0.1	0.1
c_1	Baseline production cost per barrel	\$10	\$10
c_2	Marginal cost of scaling production	\$0.1	\$0.1
ρ	Rate of Inflation	0.02	0.02/12

Table 1: Model Parameters: Descriptions and Conversions to Monthly Units

$$dN_t = \begin{cases} 1 & \text{w.p } \lambda\Delta t \\ 0 & \text{w.p } 1-\lambda\Delta t \end{cases} \tag{10}$$

The agent's objective is to maximize the expected total reward:

$$\mathcal{J}(\theta) = \mathbb{E}\left[\sum_{n=0}^{N-1} r_{t_n}\Delta t\right] \tag{11}$$

$$r_t = e^{-\rho(t-T)}(\min(q_t, D_t)P_t - C(q_t)) \tag{12}$$

With that we can outline a general training strategy to recover the optimal control q^*.

Algorithm 1 Monte Carlo Neural Network Training

Input: Number of paths per epoch M
Input: Number of training epochs T
Initialize: neural network policy $\pi_\theta(D, P, t)$
for $t = 1, \ldots, T$ **do**
  for $i = 1$ to M **do**
    Simulate trajectory $\{(D_t^i, P_t^i)\}_{t=0}^{T}$ with control $q_t^i = \pi_\theta(D_t^i, P_t^i, t)$
    Compute cumulative reward: $R^i = \sum_{n=0}^{N-1} r_{t_n}^i \Delta t$
  end for
  Compute loss:

$$\mathcal{L}(\theta) = -\mathcal{J}(\theta) = -\frac{1}{M}\sum_{i=1}^{M} R^i$$

  Update θ by gradient descent on $\mathcal{L}(\theta)$
end for

2.2.1 Results

We ran the Monte Carlo algorithm 1 for 100 epochs, with the parameters defined in table 1. We find that the supplier exhibits myopic short term value maximization behavior. Specifically the optimal strategy found is to produce no oil in the first couple of months of the simulation to create massive oil shortages. That shortage will in turn cause a jump in the price of oil. The supplier then exploits that jump in price by suddenly ramping up production to sell as much high priced oil as possible. Also note that the production tracks demand closely but is often times greater then demand. Under our model assumptions any oil produced that is not immediately sold is discarded. So while slight overproduction might seem un-intuitive it is simply the producers strategy of hedging against the volatility in demand and ensuring that they do not miss out on any potential profit.

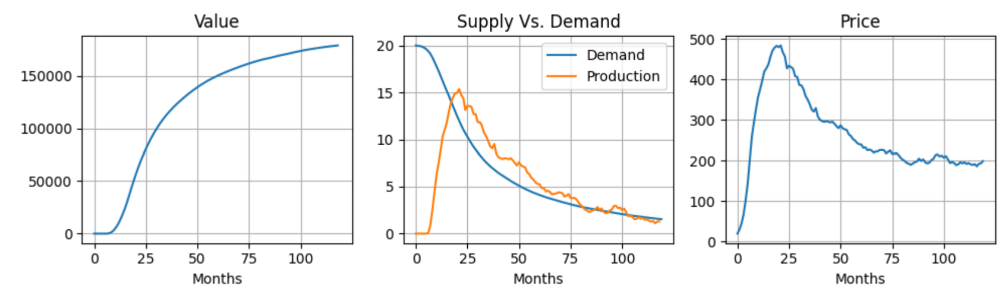

Figure 1: Demand collapsing production strategy with T equal to 10 years and dt equal to one month

This optimal production strategy is certainly valid, but is not realistic reflection of real world market dynamics since an oil producer would want to retain demand overtime. Instead demand collapses within a few time-steps. We can attribute this directly to the β and α_D parameter. Since if $\beta P_t < \alpha_D$ then the drift term of demand will consistently be negative. In other words the downward pressure of high prices always dominates the current demand for oil in the marketplace. Adjusting these parameter to $\alpha_D = 0.6$ and $\beta = 0.001$ yields a similar but more sustainable oil producing strategy, where the producer periodically creates shortages in order to induce a jump in price which the producer then exploits by increasing production. However as a result of the increase in production, a surplus is created which crashes the price and decreases demand. The producer then responds to the crash in demand by creating a shortage. This cyclical pattern of behavior, while volatile, proves to be highly profitable for the producer.

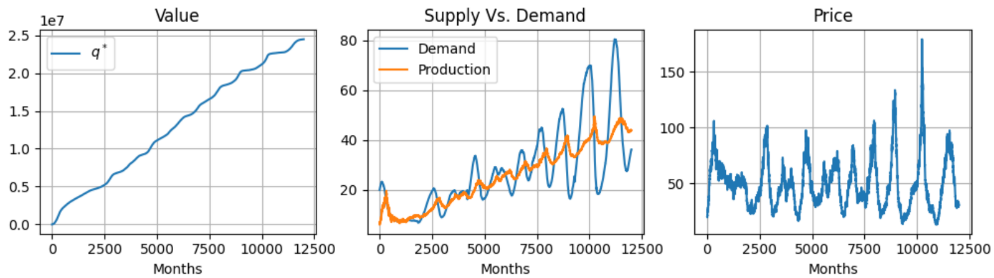

Figure 2: Oscillatory production strategy

This optimal strategy is a reasonable strategy under our simplified modeling framework. However its not reflective of real world market dynamics since such volatile behavior could catalyze recessionary market conditions and would inevitably persuade market participants to rely on less volatile energy solutions. We note that such oscillatory behavior is a natural form of price control. What we take issue with is the intensity of the swings in price. In order to dampen the swings to create a more realistic we model we adjust $\eta = 1.01$, which increases price elasticity and indicates that the price responds more conservatively to shortages and surpluses.

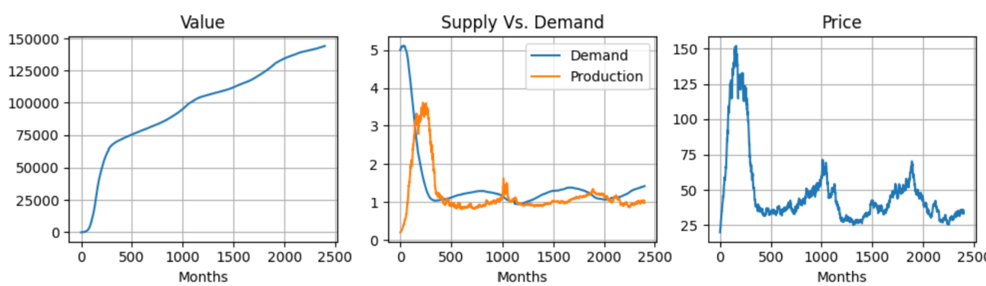

Figure 3: Stable oscillatory production strategy

An obvious production strategy is just to produce as much oil as there is demand. However by doing this drift term of price is effectively 0 which means that P hovers around P_0. So as demand increases the producer is unable to turn a profit since the cost is quadratic in q, and the price remains constant in expectation. In other words for all $t > \tau = \inf\{t : c_1 q_t + c_2 q_t^2 \geq q P_0\}$ the producer operates at a loss.

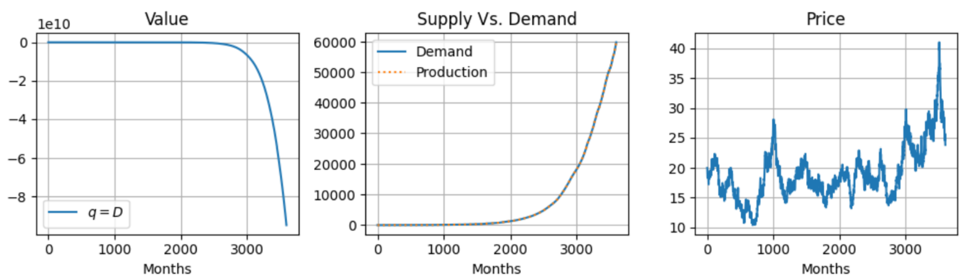

Figure 4: $D_t = q_t$ production strategy

At this point it clear that when the model uses price dynamics 13 with drift $\eta \log(D_t/q_t)$ the price is highly sensitive to small shortages and surpluses. This opens the door for producers to engage in extreme price control/manipulation in order to guarantee long term revenue. The degree to which a producer induces boom and bust cycles is theoretically unconstrained, but in reality the producer will most likely limit the magnitude of their manipulation to maintain consumer interest. To that end we introduce a variation to the model to limit the producers access to the boom-bust strategy. Specifically we let

$$dP_t = P_t[\eta \tanh(D_t - q_t)dt + \sigma_P dW_t^P] \tag{13}$$

In this case when there is a market surplus $q_t > D_t$ there is a negative drift and when there is a market shortage $q_t < D_t$ there is a positive drift in price. However unlike the log based model the drift in this case is bounded from $[-\eta, \eta]$. This alteration forces the producer to engage in more conservative strategies since their ability to manipulate the market price is significantly diminished. Under this setting the optimal strategy is to slightly but consistently underproduce. Specifically we find that across $M = 10^4$ simulations with $T = 200$

$$\frac{1}{T \cdot M} \sum_{i=1}^{M} \sum_{t=1}^{T} (D_t - q_t) \approx .4438 \ldots \tag{14}$$

This strategy is far more stable then the boom-bust strategy we observed earlier while sacrificing little to no value on average. Additionally we observe that under the tanh modeling choice price stabilizes and demand is gradually forced upward, leading to a linear increase in revenue generated over time. An added benefit of the price dynamics 29 is that the drift is Lipschitz which allows us to use Piccard iteration to solve the problem, which we focus on in the next section.

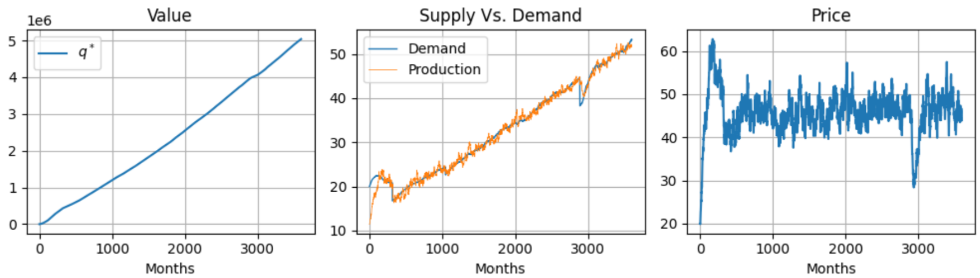

Figure 5: Optimal strategy under $\mathbb{E}[dP_t] \propto \tanh(D_t - q_t)$

2.3 Piccard Iteration Scheme

The main idea behind the finite difference scheme is to discretize the state variable (D_t, P_t, t) then use approximations of the partial derivatives to retrieve a solution. However in our case the Hamiltonian of the HJB PIDE involves a supremum that depends on $\partial_P V_{i,j}^n$, which itself depends on the unknown solution V^n . Traditional finite difference schemes assume known coefficients. In this case however the coefficients of the discretized system depend nonlinearly on the solution. As a result we can't use the traditional finite difference scheme. To that ends we choose to apply the Piccard iteration method, and begin by formalizing a discretization scheme. We note that in practice $\Delta t = (\Delta D)^2 = (\Delta P)^2$, and that we will use the price dynamics described in equation 29 in order to preserve Lipchitz continuity and guarantee convergence.

$$t_n = n\Delta t, \quad n = 0, 1, \ldots, N_t \quad \Delta t = \frac{T}{N_t}$$

$$D_i = D_{\min} + i\Delta D, \quad i = 0, 1, \ldots, N_D \quad \Delta D = \frac{D_{\max} - D_{\min}}{N_D}$$

$$P_j = P_{\min} + j\Delta P, \quad j = 0, 1, \ldots, N_P \quad \Delta P = \frac{P_{\max} - P_{\min}}{N_P}$$

Next we approximate the derivatives using finite differences. We denote the approximation of $V(D, P, t)$ by $V_{i,j}^n \approx V(D_i, P_j, t_n)$.

$$\partial_t V \approx \frac{V_{i,j}^{n+1} - V_{i,j}^n}{\Delta t} \tag{15}$$

$$\partial_D V \approx \frac{V_{i+1,j}^n - V_{i-1,j}^n}{2\Delta D} \tag{16}$$

$$\partial_P V \approx \frac{V_{i,j+1}^n - V_{i,j-1}^n}{2\Delta P} \tag{17}$$

$$\partial_{DD}^2 V \approx \frac{V_{i+1,j}^n - 2V_{i,j}^n + V_{i-1,j}^n}{(\Delta D)^2} \tag{18}$$

$$\partial_{PP}^2 V \approx \frac{V_{i,j+1}^n - 2V_{i,j}^n + V_{i,j-1}^n}{(\Delta P)^2} \tag{19}$$

$$\partial_{DP}^2 V \approx \frac{V_{i+1,j+1}^n - V_{i-1,j+1}^n - V_{i+1,j-1}^n + V_{i-1,j-1}^n}{4\Delta D \Delta P} \tag{20}$$

We approximate the jump term $\mathbb{E}_Y[V(D_i e^Y, P_j, t_n)]$ using Guass-Hermite quadrature since $Y \sim \mathcal{N}(\mu, \delta^2)$. Where $(z_k, w_k)_{k=1}^M$ are Gauss-Hermite Quadrature points and weights of order M. Note that $D_i e^Y$ may not perfectly line up with a node on our D grid. So for any y we find a node D_m such that $D_m \leq D_i e^y \leq D_{m+1}$ and interpolate.

$$\int_{-\infty}^{\infty} V(D_t e^y, P_t, t)\phi(y)dy \approx \sum_{k=1}^{M} \hat{w}_k \tilde{V}(D_i e^{y_k}, P_j, t_n) \tag{21}$$

$$y_k = \sqrt{2}\delta + \mu, \quad \hat{w}_k = \frac{w_k}{\sqrt{\pi}} \tag{22}$$

$$\tilde{V}(D_i e^{y_k}, P_j, t_n) \approx \frac{D_{m+1} - D_i e^{y_k}}{D_{m+1} - D_m} V_{m,j}^n + \frac{D_i e^{y_k} - D_m}{D_{m+1} - D_m} V_{m+1,j}^n \tag{23}$$

$$\text{s.t } D_m \leq D_i e^{y_k} \leq D_{m+1} \tag{24}$$

The last component necessary to fully discretize the HJB PIDE is to approximate the policy q_{ij}^n. To do this we must maximize the Hamiltonian which forces the producer to balance immediate revenue gain with future negative price drift.

$$J_{i,j,n}(q) = e^{-\rho(t^n - T)}\left[\min(q, D_i)P_j - C(q)\right] + \eta \tanh(D_t - q_t)P_j \partial_P V_{i,j}^n \tag{25}$$

$$q_{ij}^{n*} = \arg\sup_{q \geq 0} \{J_{i,j,n}(q)\}, \quad \mathcal{H}_{i,j,n} = \sup_{q \geq 0} J_{i,j,n}(q) \tag{26}$$

Now we can put these components together and fully discretize the HJB PIDE and compactly describe the Piccard iteration method. Where $V^{(k)}(D_i, P_j, t_n) = V_{i,j,n}^{(k)}$ at the k^{th} iteration, and all partial derivatives are computed using the the finite difference scheme described above.

$$
\begin{aligned}
F_{i,j,n}^{(k)} &= (\partial_D V)_{i,j,n}^{(k)} \cdot D_i(\alpha_D - \beta P_j) + \frac{1}{2}(\partial_{DD}^2 V)_{i,j,n}^{(k)} \cdot D_i^2 \sigma_D^2 \\
&+ \frac{1}{2}(\partial_{PP}^2 V)_{i,j,n}^{(k)} \cdot P_j^2 \sigma_P^2 + (\partial_{DP}^2 V)_{i,j,n}^{(k)} \cdot D_i P_j \sigma_D \sigma_P \rho \\
&+ \lambda \left(\sum_{k=1}^{M} \hat{w}_k \tilde{V}^{(k)}(D_i e^{y_k}, P_j, t_n) - V_{i,j,n}^{(k)} - k D_i (\partial_D V)_{i,j,n}^{(k)} \right) \\
&+ \mathcal{H}_{i,j,n}^{(k)}
\end{aligned}
\tag{27}
$$

Then we can solve for V at each (i, j, n) grid point by iteratively updating $V_{i,j}^n$ as

$$
\begin{aligned}
0 &= \frac{V_{i,j}^{n+1} - V_{i,j,n}^{(k)}}{\Delta t} + F_{i,j,n}^{(k)} \\
V_{i,j,n}^{(k+1)} &\leftarrow V_{i,j}^{n+1} - \Delta t \cdot F_{i,j,n}^{(k)}
\end{aligned}
\tag{28}
$$

Algorithm 2 Picard Iteration for Solving HJB-PIDE via Finite Differences

Set maximum iterations K and tolerance δ
Set N_t, N_D, N_P and initialize grids $\{D_i\}_{i=1}^{N_D}, \{P_j\}_{j=1}^{N_P}$
Set parameters $\eta, \sigma_P, \alpha_D, \beta, \lambda, \sigma_D, \mu, \delta, c_1, c_2$
Initialize $V_{i,j}^{N_t} = 0 \quad \forall i, j$ (terminal condition)
for $n = N_t - 1, N_t - 2, \ldots, 0$ **do**
    Set initial guess $V_{i,j}^{(0)} = V_{i,j}^{n+1}$
    for $k = 0, 1, \ldots, K$ **do**
        for $(i, j) \in \{D_i\} \times \{P_j\}$ **do**
            Compute $F_{i,j,n}^{(k)}$ using $V_{i,j,n}^{(k)}$
            Update:
$$V_{i,j,n}^{(k+1)} = V_{i,j}^{n+1} - \Delta t \cdot F_{i,j,n}^{(k)}$$
        end for
        if $\max_{i,j} \left| V_{i,j,n}^{(k+1)} - V_{i,j,n}^{(k)} \right| < \delta$ **then**
            break
        end if
    end for
    Set $V_{i,j}^n = V_{i,j,n}^{(k+1)}$
end for
Output: $V_{i,j}^0$

We run this algorithm for $T = 100$ and plot the results for $V(0, D_0, P_0)$ for $D_0, P_0 \in [20, 120]$ with the same parameter in table 1 except we let $\alpha_D = .6, \beta = .001, \eta = 1.01$. We observe that the solution is similar to the optimal production generated by the Monte Carl method with tanh price dynamics as described in equation 29 and figure 5. Namely the optimal strategy is still to slightly but consistently underproduce to drive up demand and stabilize prices.

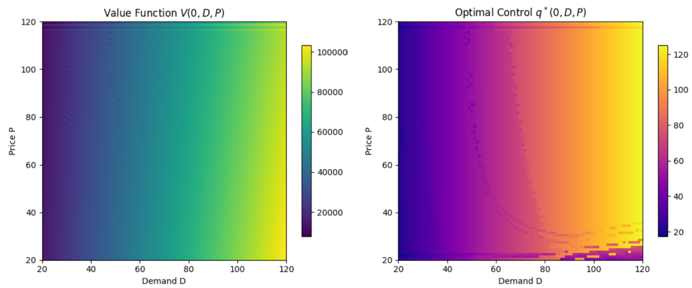

Figure 6: Optimal strategy $q(0, D_t, P_t)$ and value $V(0, D_t, P_t)$ under $\mathbb{E}[dP_t] \propto \tanh(D_t - q_t)$ using Piccard iteration

We can also plot a random path that uses the optimal strategy found by the Piccard iteration 7. Recall that for the Monte Carlo we found that the producer averaged a shortage of about .443, similarly in this case we observe that across $M = 10^4$ simulations with $T = 200$

$$\frac{1}{T \cdot M} \sum_{i=1}^{M} \sum_{t=1}^{T} (D_t - q_t) \approx .316 \ldots \tag{29}$$

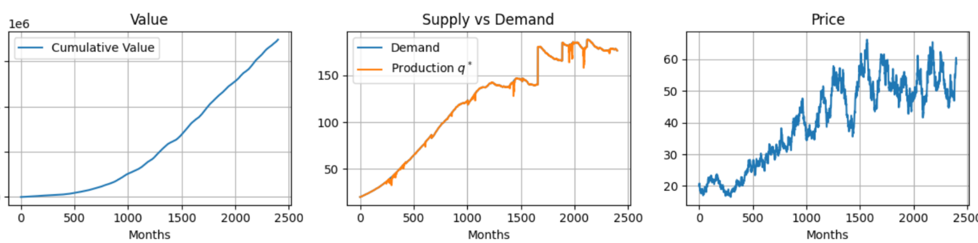

Figure 7: Example of Optimal strategy under $\mathbb{E}[dP_t] \propto \tanh(D_t - q_t)$ using Piccard iteration for random D_t, P_t Path

3 Conclusion

In this short paper, we developed a stochastic closed-loop control model for oil production under demand uncertainty and market shocks. We modeled price and demand dynamics as coupled stochastic differential equations, with both Brownian volatility and Poisson-driven jumps to reflect real-world disruptions in oil markets. We introduced a revenue focused value funciton and derived the corresponding Hamilton-Jacobi-Bellman Partial Integro-Differential Equation (HJB-PIDE).

We explored two numerical approaches to solving the HJB-PIDE, including a Monte Carlo neural network method and a Picard iteration finite difference scheme. We found that the Monte Carlo method was naturally far more flexible and could produce a wider variety of strategies since we don't need to worry about the dynamics being Lipchitz. However across these production strategies, we observed consistent trends in optimal policy behavior. Specifically, producers strategically underproduce to generate shortages, resulting in price hikes, which increase future revenues. Under the $\log(D/q)$ price dynamics, this behavior led to unstable but profitable boom-bust cycles. However, by replacing the log drift with a bounded $\tanh(D-q)$ function, the model produces more realistic, stable strategies that promote sustained demand growth and revenue.

Our results suggest that the producer can exploit the feedback loop between production and price to induce favorable market conditions. The degree of control depends heavily on the elasticity parameter η and the form of and parameters contained in the price drift.

Future work could extend the model by incorporating inventory constraints and storage costs. Additionally, modeling competition between multiple producers and volatility induced consumer drop-out, could yield a richer framework. Regardless, our current work provides a simple and tractable approach for analyzing optimal commodity production strategies in volatile and discontinuous markets.

References

[1] Robert C. Merton. Option pricing when underlying stock returns are discontinuous. *Journal of Financial Economics*, 3(1-2):125–144, 1976.

[2] Rama Cont and Peter Tankov. *Financial Modelling with Jump Processes*. Chapman & Hall/CRC Financial Mathematics Series. Chapman & Hall/CRC, Boca Raton, 2004.

[3] Dario Caldara, Michele Cavallo, and Matteo Iacoviello. Oil price elasticities and oil price fluctuations. International Finance Discussion Papers 1173, Board of Governors of the Federal Reserve System (US), 2016.

[4] U.S. Energy Information Administration. Short-term energy outlook supplement: U.s. crude oil production to 2050 in three price and technology cases. Technical report, U.S. Department of Energy, 2021. Accessed: 2025-04-28.

[5] John Asker, Allan Collard-Wexler, and Jan De Loecker. *Market Power, Production (Mis)Allocation and OPEC*. September 2017.