in Table 1 were chosen to cover a wide range of scenarios for demand and cost parameters.

For each of the 1000 trials, we applied our 3-step approach to find near-optimal values of $s, S,$ and p. We varied p from $1.1\frac{\lambda}{\mu}$ to $5.0\frac{\lambda}{\mu}$ using the same increments as in Example 1, resulting in 67 values of p considered.

Figure 5 shows the frequency of the best production rate factors for the 1000 trials. Over 80% of the trials resulted in a best production rate factor between 1.125 and 1.9 (equivalent to utilizations of 0.89 and 0.53, respectively). Some of the trials (8%) resulted in a production rate factor equal to the maximum rate of 5.0; these trials had unit production cost functions that favored high production rates.

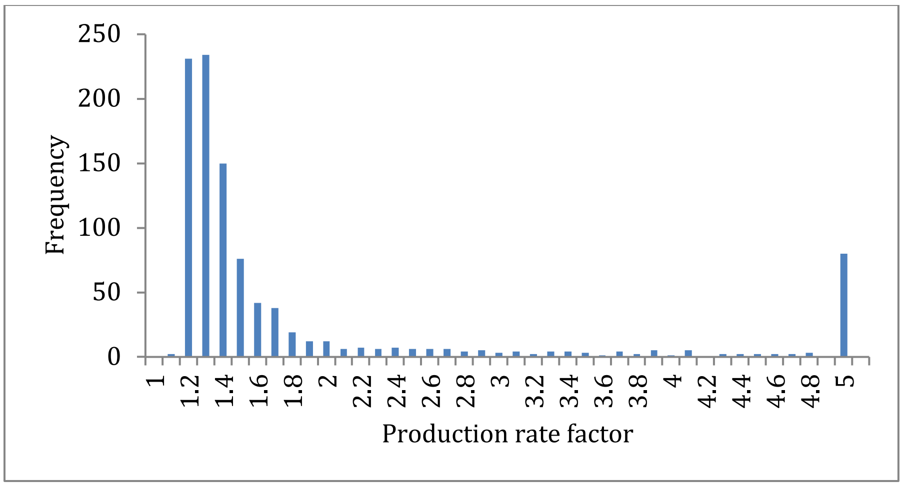

FIGURE 5: FREQUENCY OF BEST PRODUCTION RATE FACTOR, STOCHASTIC CASE

To give a partial view of the numerical study, we show results from a sample of the trials in Appendix A. The sample was selected by sorting the trials by the best production rate factor from smallest to largest and selecting the first trial in the list and every 100th trial for a total of 11 trials.

How do our results for the stochastic trials compare to that of equivalent deterministic trials? To make this comparison, for each trial in the numerical study, we created an equivalent deterministic trial by setting the constant demand D equal to the expected demand λ/μ. Then, we solved for the best production rate p via a one-dimensional search of the total cost function TC in (3) over values of p from $1.1D$ (production rate factor 1.1) to $5.0D$ (production rate factor 5.0). Figure 6 shows the frequency of the best production rate factors for these 1000 trials.

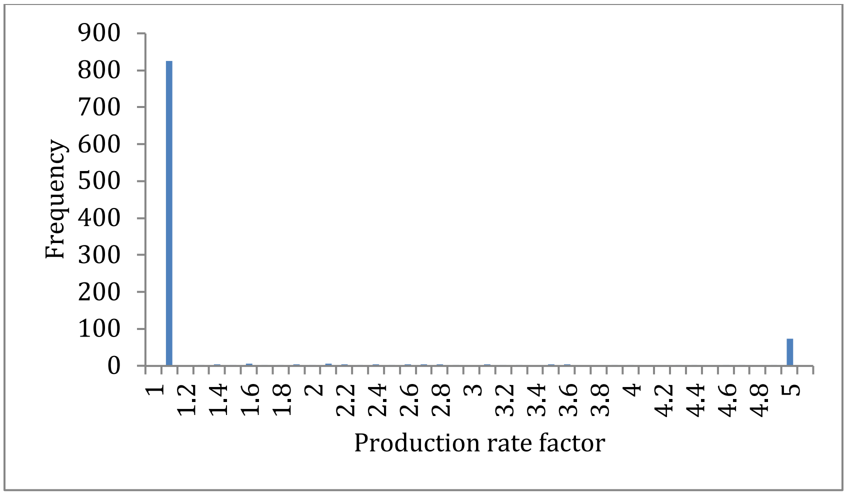

FIGURE 6: NUMBER OF TRIALS BY BEST PRODUCTION RATE FACTOR, DETERMINISTIC CASE

The deterministic case favors the minimum production rate factor of 1.1, while the stochastic case tends to result in production rate factors greater than the minimum rate. Table 2 shows the number of trials for which the best production rate factor was at the minimum of 1.1, between 1.1 and 5.0, and at the maximum of 5.0.

TABLE 2: TALLY OF TRIALS BY BEST PRODUCTION RATE FACTOR

Best production rate factor	Deterministic case	Stochastic case
Minimum of 1.1	825	2
Between 1.1 and 5.0	102	918
Maximum of 5.0	73	80

To explain these differences, we look at the setup, inventory holding, and backorder costs for the deterministic and stochastic cases. We do not need to compare the production costs because they are the same for the two cases. In the deterministic case, the average setup, inventory holding, and backorder costs per unit time each are concave increasing in *p*; thus, the sum of the three costs are concave increasing in *p*. In contrast, the expected setup, inventory holding, and backorder costs per unit of time in the stochastic problem take on different shapes. In Figs. 7 and 8, we show the relevant costs for the two cases, deterministic and stochastic, respectively, for Trial 1 ($\lambda = 86$, $\mu = 0.040$, $K = 631$, $h = 3$, $b = 42$, $V_{min} = 8.5$, $\varepsilon = 0.133$, $O_{min} = 6.2$, $\beta = 0.461$). The top curve in each figure is the sum of the three costs (setup, inventory holding, and backorder). For the deterministic case, the concave increasing shape of the sum of the three costs explains why this case favors setting *p* equal to the minimum production rate. For the best *p* to be greater than the minimum production rate, the production costs $Df(p)$ must decrease in *p* more than the setup, inventory holding, and backorder costs increase in *p*. In the stochastic case, when *p* is low $\left(\text{close to } \frac{\lambda}{\mu}\right)$, the utilization is close to 100%, and the expected backorder and inventory holding costs per unit time are very high. As with queuing systems, production systems with high utilization are expensive. This phenomenon explains why the stochastic case favors setting *p* above the minimum value of *p* = 1.1 (equivalently, setting the utilization below 0.91). Moreover, we would expect that, as demand becomes more intermittent (lower λ)

with larger demand sizes (smaller μ) (i.e., lumpier demand), the system would favor higher production rates. In the next section, we study the effect of lumpier demand on the best production rate.

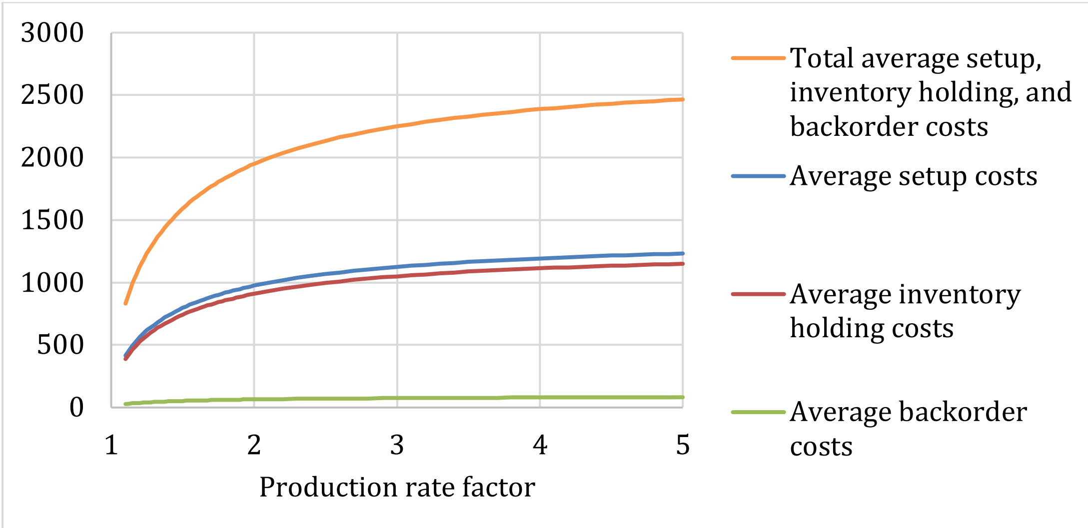

FIGURE 7: TRIAL 1: AVERAGE SETUP, INVENTORY HOLDING, AND BACKORDER COSTS IN DETERMINISTIC CASE

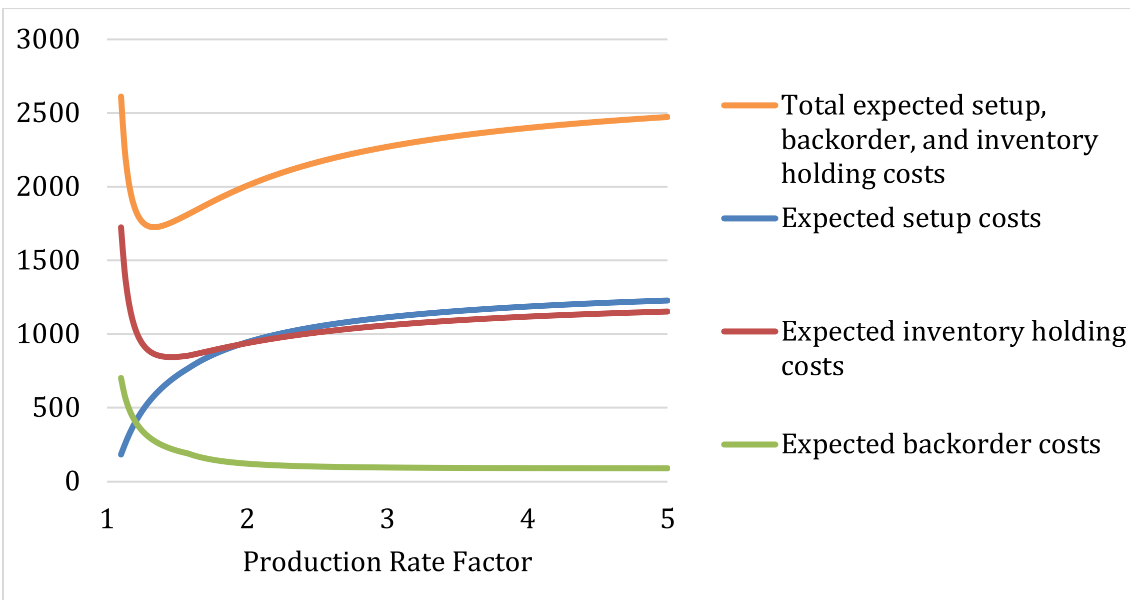

FIGURE 8: TRIAL 1: EXPECTED SETUP, INVENTORY HOLDING, AND BACKORDER COSTS IN STOCHASTIC CASE

4.2 Sensitivity Analysis on the Randomness of Demand

In this section, we examine the sensitivity of the best production rate factor to changes in the level of lumpiness of demand. Demand is considered lumpier if λ and μ are smaller for a given expected demand rate. For example, given the same expected demand λ/μ = 2500, demand with λ = 25 arrivals per time unit and μ = 0.01 (expected demand size = 100) is lumpier (i.e., more intermittent with larger demand sizes) than demand with λ = 100 arrivals per time unit and μ = 0.04 (expected demand size = 25). We note that lumpier demand has a higher variance. The variance for a compound Poisson demand process with exponential demand sizes is equal to $2\lambda/\mu^2$; thus, for a given expected demand λ/μ, the variance increases as μ decreases.

To conduct this sensitivity analysis, we expanded the numerical study in the previous section. For each of the 1000 trials in the numerical study in Section 4.1, we created five sub-trials with different λ and μ values ranging from low λ and μ values (lumpiest) to high λ and μ values (least lumpy) while keeping λ/μ fixed at the original value. We used multipliers 0.25, 0.50, 1.0, 2.0, and 3.0 of the original λ and μ values. For example, applying the multiplier of 0.25 results in one-quarter of the original values for λ and μ, while the multiplier 3.0 results in three times the original values of λ and μ. As an example, in Trial 1, the original λ and μ values were 86 and 0.04, respectively, and we created five sub-trials with λ and μ values, as shown in Table 3. All five sub-trials have the same expected demand λ/μ = 2150 and cost parameters as in the original Trial 1.

TABLE 3: EXAMPLE: SUB-TRIALS FOR TRIAL 1

	Sub-trial 1 (lumpiest)	Sub-trial 2	Sub-trial 3 (Original values)	Sub-trial 4	Sub-trial 5 (least lumpy)
Multiplier of original values for λ and μ	0.25	0.5	1.0	2.0	3.0
λ	21.5	43	86	172	258
μ	0.01	0.02	0.04	0.08	0.12

By repeating this approach for all 1000 trials, we ended up with 1000 sets of five sub-trials where we varied the lumpiness of demand. For each sub-trial, we identified the best production rate p. We found that in 925 of the 1000 trials, the best production rate decreased as the multiplier increased from 0.25 (lumpiest) to 3.0 (least lumpy). In 75 trials, the best production rate stayed at the maximum of 5.0 as the multiplier increased due to unit production cost functions that favored high production rates. Figure 9 shows four trials illustrating how the best production rate changes as demand gets less lumpy. See Appendix B for parameter settings for the trials in Figure 9. In addition, we found that as the multiplier increased from 0.25 (lumpiest) to 3.0 (least lumpy), the best production rate moved closer to that of the best production rate in the equivalent deterministic trial where constant demand $D = \lambda/\mu$. More frequent arrivals (larger λ) with smaller demand sizes (larger μ) result in more predictable demand, and the system favors lower production rates (i.e., higher utilization), moving toward that of the deterministic setting.

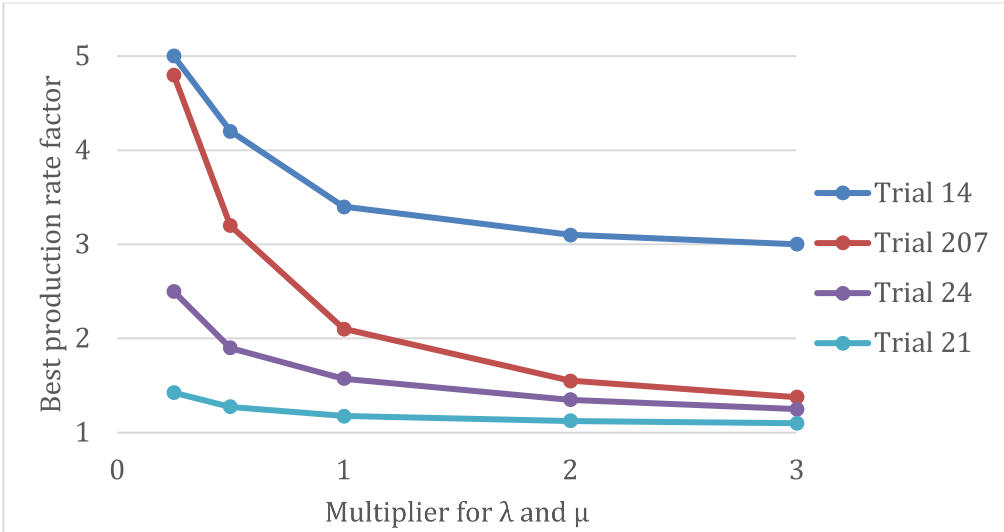

FIGURE 9: BEST PRODUCTION RATE FACTOR VERSUS MULTIPLIER FOR λ AND μ.

4.3 Sensitivity Analysis on the Cost Parameters

In this section, we examine the sensitivity of the best production rate factor to changes in the cost parameters, K, h, and b, for both the deterministic and stochastic cases. To do this, we expanded the numerical study in the previous section to vary K, h, and b one at a time. For each of the 1000 trials in the numerical study in Section 4.1, we created five sub-trials with different parameter values ranging from low to high. We used multipliers 0.25, 0.50, 1.0, 2.0, and 3.0 of the parameter's original value. As an example, in Trial 1, the original K value was 631, and to vary K, we created five sub-trials with the following K values: 158, 316, 631, 1262, 1893, while keeping all the other parameters the same as in the original numerical study. Thus, we ended up with 1000 sets of five sub-trials where we varied K. We did the same for h and b. For each sub-trial, we identified the best production rate for the deterministic and stochastic cases. We report the results in Tables 4 and 5 for the deterministic and stochastic cases, respectively, tracking the best production factor as the parameter changes from its lowest to highest value. "Not changing" means that, as the parameter value increased, the best production rate stayed the same at the lowest parameter value (at the multiplier of 0.25). "Increasing" ("Decreasing") means that the best production rate increased (decreased) as the multiplier for the parameter increased from 0.25 to 3.0.

TABLE 4: RESULTS FOR DETERMINISTIC NUMERICAL STUDIES SENSITIVITY ANALYSIS

As each parameter increases, the number of trials for which the best production rate was	K	h	b
Not changing	756	766	897
Decreasing	244	234	103

TABLE 5: RESULTS FOR STOCHASTIC NUMERICAL STUDIES SENSITIVITY ANALYSIS

As each parameter increases, the number of trials for which the best production rate was	K	h	b
Not changing	59	79	89
Decreasing	941	220	10
Increasing	0	695	862
Decreasing, then increasing	0	6	39

For the deterministic case, the best production rate is either not changing or decreasing in K, h, or b. (See Table 4.) In most of the deterministic trials, the best production rate did not change in K, h, or b. In most of these “not changing” trials, the best production rate factor at the lowest parameter value (multiplier of 0.25) was the minimum rate of 1.1, and it did not change as the multiplier increased. It is easy to show analytically that the mixed partials of the total cost function TC in (3), that is, $\frac{\partial TC}{\partial p \partial K}, \frac{\partial TC}{\partial p \partial h}, and\ \frac{\partial TC}{\partial p \partial b}$ are greater than zero, implying that the best production rate p is decreasing or not changing as K, h, or b increase. We note that Khouja (1995) and Bhandari and Sharma (1999) show examples of varying K and h that are consistent with these results.

For the stochastic case, we can see from Table 5 that, in all the trials, as K increases, the best production rate factor either decreases or does not change, similar to that of the deterministic case. However, as h or b increases, the best production rate increases in most trials. This result can be explained by looking at the shape of the three expected cost functions: setup, inventory holding, and backorder. Referring back to Fig 8, we can see that the expected setup cost takes on a concave increasing shape similar to that of the deterministic case; however, the expected inventory holding and expected backorder costs take on very different shapes than those of the deterministic case. The expected inventory holding cost is very high at low production rates, and as the production rate increases, it decreases in a convex way and then increases in a concave way. The expected backorder cost is also very high at low production rates, and as the production rate increases, it decreases in a convex way. Thus, as h or b increase, the stochastic nature of the problem makes it marginally more expensive to move from a relatively low production rate to a lower one.

To give a partial view of the results of this sensitivity study, we show in Appendix C graphs of how the best production rate changes as the parameters K, h, or b change, respectively. In each chart, we show four different examples to illustrate the different patterns of changes in the best production rate.

V. CONCLUSION

We developed a framework for a production-inventory problem with stochastic demand to identify the best production rate in addition to the two-critical-number (*s, S*) policy. The proposed framework builds upon the heuristic in Azoury and Miyaoka (2020), which gives closed-form near-optimal solutions for the two-critical-number (*s, S*) policy. These closed-form solutions are important because they vastly simplify the search for the best p, s, and S by reducing a multi-dimensional search to a one-dimensional one. The framework we proposed, which we call the 3-step approach, starts with a p, computes (*s, S*) from the closed-form solutions, plugs into the cost function, and moves on to the next p to find the best p. This one-dimensional search over p values can easily be done in a spreadsheet.

Our results differ from prior research in a deterministic EPQ setting where the best production rate tends to be an extreme point solution. In our stochastic setting, we found through extensive numerical studies that nearly all of the best production rates were in the middle range of production values considered. Due to high inventory and backorder costs, it was rare to have the best production rate near the expected demand rate (i.e., high utilization) in our stochastic setting. In addition, we found that as the degree of randomness increases, the best production rate increases, diverging further from the deterministic case. We also studied the sensitivity of the best production rate to changes in K, h, or b. In the deterministic setting, the best production rate is decreasing or not changing in K, h, or b. Moreover, the deterministic case favors the minimum production rate. In the stochastic setting, we observed that the best production rate was decreasing or not changing in K but was increasing, decreasing, or not changing in h or b.

A direction for future research is to study a more general production-inventory model where demand is a mixture of two components: one stochastic and one deterministic. For a fixed production rate, Azoury and Miyaoka (2020) studied this model with a two-critical-number policy and derived expected cost functions. They proposed closed-form, near-optimal solutions for the two-critical-number policy. Adding a deterministic component to demand is a significant addition because transitions in the system occur continuously and at demand arrival epochs. It would be interesting to extend the framework we developed here with the production rate as a decision variable to this more general case where demand is a mixture of stochastic and deterministic components.

REFERENCES

Afshar-Nadjafi, B., "An EPQ model with unit production cost and set-up cost as functions of production rate." *Modelling and Simulation in Engineering*, 2013, https://doi.org/10.1155/2013/727685.

Azoury, K. S., and Miyaoka, J., "Optimal and simple approximate solutions to a production-inventory system with stochastic and deterministic demand." *European Journal of Operational Research*, 286(1), 2020, 178-189.

Bhandari, R. M., and Sharma, P. K., "The economic production lot-size model with variable cost function." *OPSEARCH*, 36(2), 1999, 137-150.

Chang, J., Lu, H., and Shi, J., "Stockout risk of production-inventory systems with compound Poisson demands." *Omega* 83, 2019, 181-198.

Cheng, T. C. E., "A note on an economic production quantity model and just-in-time-production." *Production Planning & Control*, 4(1), 1993, 88-92.

Gallego, G., "IEOR 4000: Production Management, Lecture 2.", lecture notes, Columbia University, New York, 2004.

Gavish, B., and Graves, S. C., "Technical note: A one-product production/inventory problem under continuous review policy." *Operations Research*, 28(5), 1980, 1228-1236.

Graves, S. C., and Keilson, J., "The compensation method applied to a one-product production/inventory problem." *Mathematics of Operations Research*, 6(2), 1981, 246-262.

Katehakis, M. N., Melamed, B., and Shi, J., "Optimal replenishment rate for inventory systems with compound Poisson demands and lost sales: a direct treatment of time-average cost." *Annals of Operations Research*, 2015, https://doi.org/10.1007/s10479-015-1998-y.

Khouja, M., "The economic production lot size model under volume flexibility." *Computers & Operations Research*, 22(5), 1995, 515-523.

Larsen, C., " Using a variable production rate as a response mechanism in the economic production lot size model." *Journal of the Operational Research Society*, 48(1), 1997, 97-99.

Shi, J., Katehakis, M. N., Melamed, B. and Xia, Y., "Production-inventory systems with lost sales and compound Poisson demands." *Operations Research*, 62(5), 2014, 1048-1063.

Shi, J., "Optimal continuous production-inventory systems subject to stockout risk." *Annals of Operations Research*, 2016, https://doi.org/10.1007/s10479-016-2339-5.

APPENDIX A

TABLE A1: SAMPLE OF NUMERICAL STUDY FOR STOCHASTIC CASE

	Parameters								Results		
Trial	λ	μ	K	b	$V(p_{min})$	ε	$O(p_{min})$	β	best production rate factor	(s, S)	utilization
500	100	0.080	975	34	8.3	0.114	8.5	0.566	1.100	(171.8, 601)	0.91
263	65	0.022	812	26	6.1	0.169	6.1	0.699	1.175	(320.3, 1154)	0.85
565	55	0.083	903	31	9.7	0.120	5.5	0.369	1.200	(50.4, 378)	0.83
993	80	0.058	605	27	9.1	0.317	5.2	0.569	1.225	(55, 464)	0.82
588	74	0.069	787	35	7.8	0.693	7.2	0.594	1.275	(30.9, 438)	0.78
586	82	0.030	554	37	7.7	0.540	5.7	0.619	1.325	(100.9, 718)	0.75
356	76	0.077	735	32	9.1	0.557	6.5	0.500	1.400	(-3.4, 402)	0.71
716	38	0.074	711	29	7.3	0.510	6.0	0.338	1.525	(-5.1, 310)	0.66
758	89	0.059	234	43	9.6	0.269	7.8	0.229	1.875	(-4.5, 345)	0.53
602	11	0.054	249	30	9.8	0.532	5.4	0.312	3.900	(-9.5, 147)	0.26
873	52	0.030	224	47	8.1	0.577	7.4	0.144	5.000	(-20.5, 426)	0.20

Note: $h = 3$ for all trials.

APPENDIX B

Table B1 shows the parameter settings for the trials in Figure 9. The λ and μ values listed in the table are associated with the sub-trial with multiplier = 1.0. To get λ and μ values for another sub-trial, multiply the λ and μ values in the table by the appropriate multiplier (0.25, 0.50, 2.0, or 3.0).

TABLE B1: PARAMETER SETTINGS FOR TRIALS IN FIGURE 9

Trial	λ at 1.0 multiplier	μ at 1.0 multiplier	K	h	b	$V(p_{min})$	ε	$O(p_{min})$	β
14	12	0.036	495	3	37	8.4	0.635	6.9	0.256
207	14	0.057	359	3	47	7.3	0.322	7.5	0.226
24	21	0.026	584	3	47	9.5	0.256	6.7	0.396
21	51	0.072	857	3	47	5.1	0.320	7.4	0.619

APPENDIX C

Figure C1 shows four trials, each with five sub-trials, to illustrate how the best production rate changes as the parameter K changes. Table C1 shows the parameter settings for the trials in Figure C1. The K values in the table are associated with the sub-trial with multiplier = 1.0. To get the K value for another sub-trial, multiply the K value in the table by the appropriate multiplier (0.25, 0.50, 2.0, or 3.0).

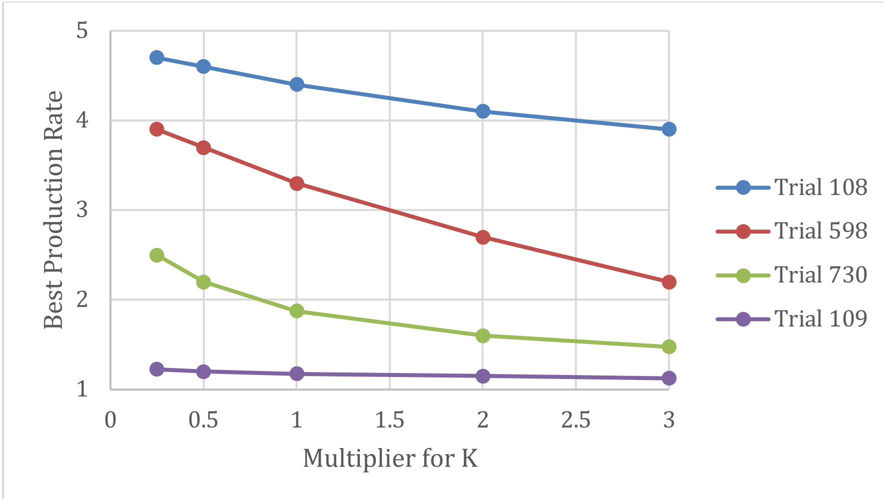

FIGURE C1: BEST PRODUCTION RATE FACTOR VERSUS MULTIPLIER FOR *K*

TABLE C1: PARAMETER SETTINGS FOR TRIALS IN FIGURE C1

Trial	λ	μ	K at multiplier 1.0	h	b	$V(p_{\min})$	ε	$O(p_{\min})$	β
108	96	0.089	643	3	21	9.8	0.556	6.2	0.259
598	25	0.026	968	3	46	8.7	0.477	7.2	0.230
730	19	0.032	689	3	45	8.1	0.257	8.9	0.184
109	69	0.099	446	3	22	6.1	0.363	9.2	0.372

Figure C2 shows four trials, each with five sub-trials, to illustrate how the best production rate changes as the parameter h changes. Table C2 shows the parameter settings for the trials in Figure C2. The h values in the table are associated with the sub-trial with multiplier = 1.0. To get the h value for another sub-trial, multiply the h value in the table by the appropriate multiplier (0.25, 0.50, 2.0, or 3.0).

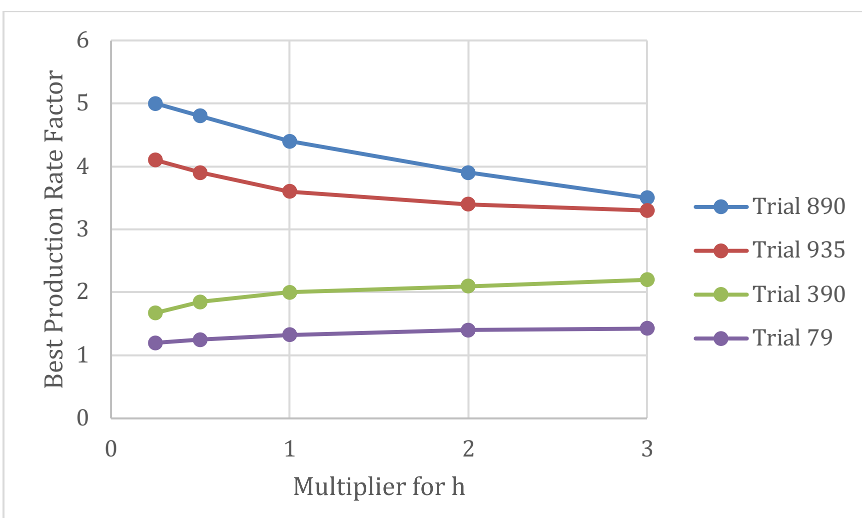

FIGURE C2: BEST PRODUCTION RATE FACTOR VERSUS MULTIPLIER FOR *h*

TABLE C2: PARAMETER SETTINGS FOR TRIALS IN FIGURE C2

Trial	λ	μ	K	h at multiplier 1.0	b	$V(p_{\min})$	ε	$O(p_{\min})$	β
890	36	0.079	696	3	41	8.0	0.620	8.7	0.160
935	13	0.024	852	3	42	7.1	0.481	10.0	0.141
390	20	0.031	203	3	22	8.6	0.110	8.6	0.113
79	33	0.042	436	3	38	5.4	0.162	6.6	0.359

Figure C3 shows four trials, each with five sub-trials, to illustrate how the best production rate changes as the parameter b changes. Table C3 shows the parameter settings for the trials in Figure C3. The b values in the table are associated with the sub-trial with multiplier = 1.0. To get the b value for another sub-trial, multiply the b value in the table by the appropriate multiplier (0.25, 0.50, 2.0, or 3.0).

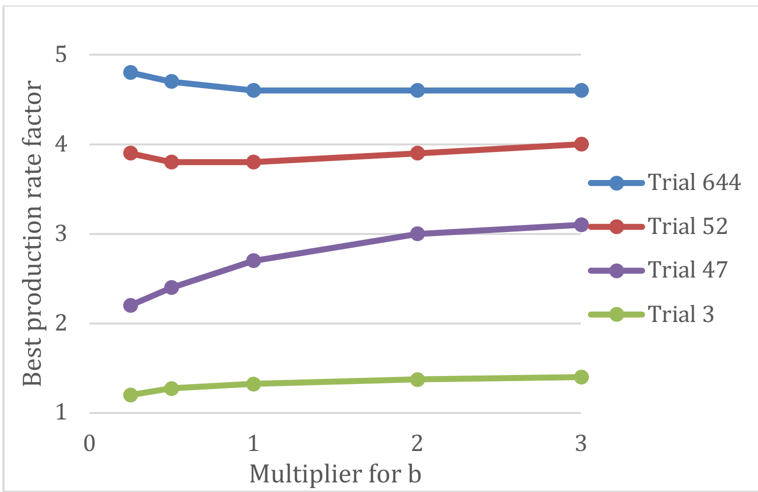

FIGURE C3: BEST PRODUCTION RATE FACTOR VERSUS MULTIPLIER FOR *b*

TABLE C3: PARAMETER SETTINGS FOR TRIALS IN FIGURE C3

Trial	λ	μ	K	h	b at multiplier 1.0	$V(p_{\min})$	ε	$O(p_{\min})$	β
644	94	0.080	981	3	20	7.9	0.580	7.0	0.189
52	94	0.091	229	3	41	9.0	0.201	7.1	0.147
47	67	0.067	754	3	28	7.2	0.489	9.6	0.653
3	40	0.071	903	3	29	5.7	0.625	5.8	0.438