Operations Management Education Review 1: 105-122.
© 2005 NeilsonJournals Publishing.

The Production Dice Game: An Active Learning Classroom Exercise and Spreadsheet Simulation

Elisabeth J. Umble and Michael Umble
Baylor University

Abstract. The basic production dice game is a powerful active learning exercise that has been extensively used in both undergraduate and graduate operations management classes as well as in numerous executive seminars and workshops. This paper goes beyond the basic model and demonstrates how eight additional models can be utilized to clearly illustrate the impact of dependency and variability on output and work-in-process inventory in balanced processes. Four additional models demonstrate the impact of dependency and variability in unbalanced processes and the importance of developing appropriate buffer and material release policies. This paper also illustrates how spreadsheet simulation can be used to reinforce the results generated in the classroom. The simulation results yield estimates of output and inventory for each model and clearly show the impact of dependency and variability on system performance.

Keywords: production dice game, spreadsheet simulation, active learning, dependency, variability, balanced plants, unbalanced plants.

1. Introduction

Some students have difficulty relating to, and maintaining a high level of interest in, Operations Management courses. Lecture-based classes exacerbate the problem of generating interesting and valuable learning experiences for the students. Garvin (1991) suggests that in the traditional lecture mode of classroom instruction, as much as fifty percent of the lecture material is forgotten within a few months. As emphasized by Heineke and Meile (1995), learning is best (and the most fun) when the learner is actively engaged with the material. Thus, a good alternative to traditional classroom lectures is active learning where the student actively participates in a structured exercise, game, or simulation designed to develop understanding of key concepts.

This paper describes how a stimulating hands-on classroom simulation known as the production dice game can be modified to illustrate several important operations management principles. The production dice game was introduced in the first edition of *The Goal* by Goldratt (1984) and more fully explained by Umble and Srikanth (1990). This paper expands upon these previous works to demonstrate how the production dice game can be modified to illustrate and specify (1) the impact of dependency and variability on throughput and work-in-process inventory in balanced plants, and (2) key managerial issues in unbalanced plants.

Ragsdale (2001) has reported that students respond favorably to the use of spreadsheet-based simulations. More recently, Johnson and Drougas (2002) demonstrated how the basic production dice game can be effectively utilized to provide an Excel-based introduction to the topic of simulation. This paper illustrates how Excel spreadsheet simulation analysis can also be utilized to validate the general relationships observed during the manual production dice game simulation, derive valid statistical estimates of key parameters, and generate sufficient statistical data to conduct significance tests between the parameters of different dice game models. Appendix A describes how the authors sequence and utilize this material in a one-semester undergraduate core course in Operations Management.

2. The Basic Production Dice Game Setup

The basic production dice game models a simple production process where material is processed sequentially through several workstations. Depending on the number of students in the class, several independent production lines of between four and eight workstations should be formed. Each student usually runs one workstation. However, if needed, one or more students may be used to run a second workstation. The basic production dice game used in this paper follows the structure originally described in Umble and Srikanth (1990). The authors decided not to use the production dice game model described by Goldratt in the *Goal* (1984) because that model has no starting work-in-process inventory. A basic setup with six workstations is illustrated in Figure 1. As also shown in Figure 1 below, except for the first workstation, each workstation maintains work-in-process inventory. The first workstation takes material from raw material stores, processes the material, and passes it to the work-in-process inventory storage area for station two. Workshop station two eventually processes and moves the material to station three, etc. When a unit of material has been processed by the last workstation, it becomes system output.

Figure 1: The Basic Production Game Setup

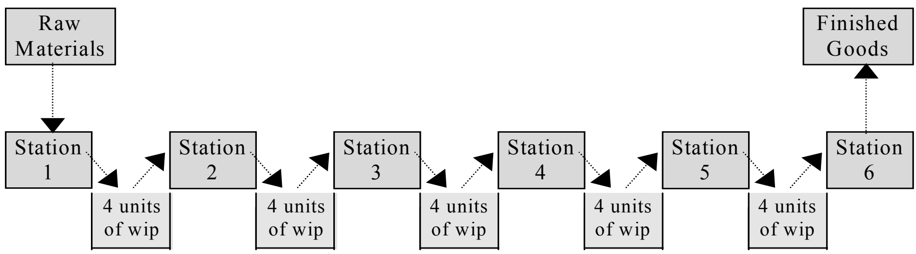

In the basic game, the roll of an ordinary gaming die is used to simulate actual daily production capacity of each individual workstation. Each simulated day consists of one toss of the die by each worker. The potential daily capacity varies from one to six units, with an average of 3.5 units. Each worker is allowed to process (move) the number of pieces determined by the roll of the die – subject to the availability of work-in-process inventory at that station at the beginning of the day. It is not permitted to use materials that were not available at the station at the beginning of the day – those units become part of the next day's work queue. Thus, it often happens that an individual workstation will not be able to produce to its daily capability due to a lack of available materials.

In the basic game, the initial starting condition is for each of the workstations to begin with four units of work-in-process inventory. This equates to a little more than one day's worth of average production capacity. The first workstation has sufficient raw materials to last for the entire simulation period.

3. Impact of Dependency and Variability on Balanced Plants

After explaining the basic dice game setup to the class (before the simulation game is performed), ask the class to determine (1) the average number of units of output a plant should be able to produce each day, and (2) the expected output after twenty days of operation. The students should answer 3.5 units and 70 units, respectively. Then ask whether they think the plant will actually achieve the theoretically expected results. Ask them to defend their answer. After the students have expressed their opinions, introduce the concepts of dependency and variability, which characterize virtually all processes.

3.1. Defining Dependency and Variability

According to Srikanth and Umble (1997), *dependency* is said to exist when certain operations or activities cannot begin until certain other operations or activities have been completed; *variability* is manifested in the form of random events and statistical fluctuations. Random events are those events that occur at irregular intervals and have a disruptive effect on the process. Statistical fluctuation refers to the idea that all processes are characterized by some degree of inherent variability. Dependency is manifested in the basic dice game by the requirement that units of material cannot be processed by a workstation until first being processed by all previous workstations. Variability is manifested by the different numbers that may occur when the dice are rolled. (It is useful to have the students identify examples of dependency and variability in real life processes.)

3.2. Balanced Plant Model Descriptions

The proposition should be presented to the class that dependency and variability will combine to degrade overall plant performance. Then a series of balanced plant models should be set up to test the hypotheses that increasing (decreasing) levels of dependency and variability will increasingly degrade (improve) performance. Balanced plant models require that every workstation have identical capacity. To provide a fair comparison between models, every workstation in each of the balanced plant models discussed here will have an average capacity of 3.5 units.

The nine models described below are only a sample of the possible models that can be used to demonstrate the desired concepts. Once each model has been simulated at least one time, the results can be compared.

Model 1: The base model (six dice). As described above, the basic dice game represents a balanced plant where six workers each roll one die to determine daily capacity. The mean absolute deviation (MAD) of each workstation is 1.5 and σ is 1.71.

Model 2: Less variability than base model (six coins; heads = 3, tails = 4). This model is the same as the base model except that the workers toss a coin instead of rolling a die. If the coin turns up "heads," then capacity for that day is three units. If the coin turns up "tails," then capacity for the day is four units. Variability in this model is much less than the base model – MAD and σ are both .5. Thus, we should expect better system performance (more output and less WIP) than in the base model.

Model 3: Similar variability as base model (six coins; heads = 2, tails = 5). This model is similar to model 2. The difference is that when the coin comes up heads, the capacity is two units. When a tail occurs, the capacity is five units. The MAD is 1.5, the same as the base model. But the σ of 1.5 is slightly less than the σ of 1.71 for the base model. Thus, system performance should be slightly better than for the base model.

Model 4: More variability than base model (six coins; heads = 1, tails = 6). This model is similar to model 2. But when the coin comes up heads, capacity is one unit. When a tail occurs, capacity is six units. The MAD and σ are 2.5, greater than the base model MAD and σ. Thus, system performance should be worse than for the base model.

Model 5: Maximum variability (six coins; heads = 0, tails = 7). This model is similar to model 2. But when the coin comes up heads, capacity is zero units. When a tail occurs, capacity is seven units. The MAD and σ are 3.5, the largest possible MAD and σ. Thus, of all the models, system performance should the worst for this model.

Model 6: Minimum dependency (two dice). This model is the same as the base model except there are only *two* stations – the lowest possible level of process dependency without eliminating dependency altogether by having only one station. Thus, system performance for this model should be better than base model performance.

Model 7: Less dependency than base model (four dice). This model is the same as the base model except there are four stations – a lower level of process dependency than the base model. Thus, system performance for this model should be better than base model performance, but worse than system performance for model 6.

Model 8: More dependency than base model (eight dice). This model is the same as the base model except there are *eight* stations – a higher level of process dependency than the base model. Thus, system performance for this model should be worse than base model performance.

Model 9: Maximum dependency (ten dice). This model is the same as the base model except there are *ten* stations – the highest level of process dependency of all the models. Thus, this model should generate the worst performance of models one and six through nine.

3.3. A Spreadsheet Analysis

Once the students have performed a "hands-on" simulation using dice or coins, then the instructor can, depending on the class, either present the spreadsheet simulation or assign it as a spreadsheet problem. The authors developed spreadsheet computer simulations for each of the above models to estimate the parameters of system performance for output and work-in-process inventory. For each model, 1000 simulation runs of 20 days each were generated.

The impact of variability. Results for models one through five illustrating the impact of variability are presented in Table 1 for the mean and standard deviation of output and ending system work-in-process inventory, starting level of inventory, individual station ending work-in-process inventory, average increase in system inventory, and average output shortfall. Recall that each workstation begins with four units of work-in-process inventory. Thus, the system begins with an initial inventory of 20 units. The average increase in system inventory is the difference between the ending level of inventory and the starting level of inventory. The average output shortfall is the difference between the theoretical expected capacity of 70 units and the actual output.

Table 1: Results of 1000 Simulation Runs to Illustrate the Impact of Variability

	Model 2 (coin 3,4) (σ = .5)	**Model 3 (coin 2,5) (σ = 1.5)**	**Model 1 (6 dice) (σ = 1.7)**	**Model 4 (coin 1,6) (σ = 2.5)**	**Model 5 (coin 0,7) (σ = 3.5)**
Mean Output	66.73	56.72	54.07	46.76	36.44
Output Standard Error	(.041)	(.115)	(.134)	(.187)	(.254)
Starting WIP	20.00	20.00	20.00	20.00	20.00
Ending WIP	23.26	33.09	35.94	43.29	53.76
Ending WIP Std. Error	(.084)	(.222)	(.259)	(.376)	(.528)
Station 2 Ending WIP	5.77	9.29	10.83	13.89	18.65
Station 3 Ending WIP	4.70	7.40	7.94	9.41	11.60
Station 4 Ending WIP	4.55	6.19	6.53	7.77	9.32
Station 5 Ending WIP	4.19	5.37	5.54	6.49	7.61
Station 6 Ending WIP	4.05	4.84	5.10	5.73	6.59
Output Shortfall	3.27	13.28	15.93	23.24	33.56
Increase In WIP	3.26	13.09	15.94	23.29	33.76

Model 1 generated an average of 54.07 units of output, and WIP increased from the starting level of 20 units to 35.94 units. Note that, as is true for each model, the increase in WIP of 15.94 units is approximately the amount of the output shortfall, 15.93. Since the first station will release approximately 70 units into the process in 20 days, all units actually released into the process that are not converted into throughput must show up as an increase in work in process.

Model 2 (with the least variability) yielded the best performance of the six models, with average output of 66.73 units and average WIP of only 23.26 units, and experienced an average output shortfall of 3.27 units. Model 3, (with similar, but slightly less, variability than the base model) yielded slightly better results than the base model, with average output of 56.72 units, WIP of 33.09 units, and an average output shortfall of 13.28 units. Model 4 (with increased variability) generated average output of 46.76 units, average WIP of 43.29 units, and average output shortfall of 23.24 units. Model 5 (with the greatest possible variability) yielded the worst performance of the six models, with average output of 36.44 units, average WIP of 53.76 units, and a whopping average output shortfall of 33.56 units.

Figure 2 presents a chart showing the standard deviation and the resulting output for each of models one through five. Interestingly enough, the data indicate a strong proportional linear relationship over the whole range of standard deviation values, from the highest level of output of 66.73 units at the minimum possible σ of .5 to the lowest level of output of 36.44 units at the maximum possible σ of 3.5.[1]

Figure 2: Output as a Function of Work Station Variablility (Standard Deviation)

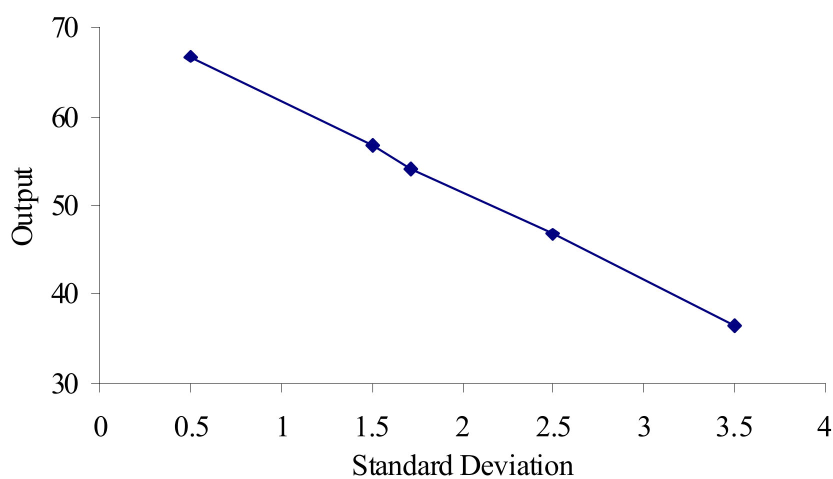

1. Baker, Powell, and Pike (1990) cite previous research indicating that for lines of different lengths, with differing processing time distributions, output decreases with increasing variability. For example, Muth (1977) showed that with three identical stations, throughput is a linear function of the coefficient of variation. Conway, Maxwell, McClain, and Thomas (1988) cite numerous previous papers showing that the loss of output in a serial line occurs in the first few stations and is a function of the variability of processing time.

Table 2: Results of 1000 Simulation Runs to Illustrate the Impact of Dependency

	Model 6 (2 dice)	Model 7 (4 dice)	Model 1 (6 dice)	Model 8 (8 dice)	Model 9 (10 dice)
Mean Output	63.21	56.78	54.07	52.95	52.63
Output Standard Error	(.200)	(.159)	(.134)	(.118)	(.102)
Starting WIP	4.00	12.00	20.00	28.00	36.00
Ending WIP	10.90	25.21	35.94	45.04	53.36
Ending WIP Std. Error	(.213)	(.258)	(.259)	(.270)	(.236)
Station 2 Ending WIP	10.90	10.71	10.83	10.73	10.91
Station 3 Ending WIP	—	8.02	7.94	7.86	7.25
Station 4 Ending WIP	—	6.48	6.53	6.47	6.70
Station 5 Ending WIP	—	—	5.54	5.71	5.85
Station 6 Ending WIP	—	—	5.10	5.19	5.08
Station 7 Ending WIP	—	—	—	4.69	4.47
Station 8 Ending WIP	—	—	—	4.39	4.64
Station 9 Ending WIP	—	—	—	—	4.34
Station 10 Ending WIP	—	—	—	—	4.12
Output Shortfall	6.79	13.22	15.93	17.05	17.37
Increase In WIP	6.90	13.21	15.94	17.04	17.36

The impact of dependency. Table 2 presents the results of 1,000 simulation runs for Models 1, 6, 7, 8, and 9. Not shown in the table is the case representing the situation where there was only one station (with resulting dependency of zero). Clearly, in such a case, the actual resulting output would be expected to average 70 units. The data indicate that increasing dependency by increasing the number of stations to two, four, six, eight and ten decreased output to 63.21, 56.78, 54.07, 52.95, and 52.63 units, respectively. As indicated in Figure 3, there is a curvilinear relationship between the number of stations and the output. This means that, initially, output losses are great with the addition of each new station. But the marginal loss of output diminishes as new stations are added.[2]

2. Several authors have shown that production line output decreases as the length of the production line grows (Gershwin, 1987). Conway, Maxwell, McClain, and Thomas (1988) show that as the number of stations increases, line output decreases, but not linearly. Multistation line output may only be 75% of the output of a single station, but the loss in output occurs mainly in the first five workstations.

Figure 3: Output as a Function of Number of Dice

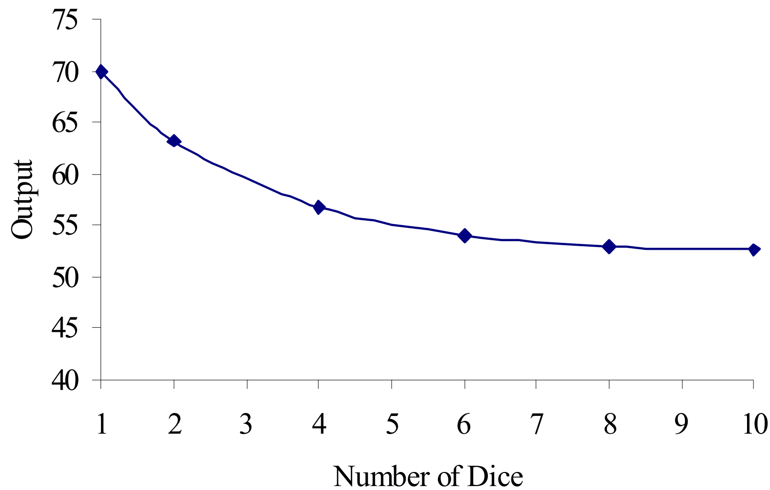

3.4. Points for Class Discussion

After the simulations have been conducted, several questions can be posed to the class to generate further discussion.

1. What would be the expected effect if models 2 and 4 were combined?

2. What would be the expected effect if models 3 and 5 were combined?

3. What can management do to reduce variability in real life processes?

4. What can management do to reduce dependency in processes?

4. Investigating Unbalanced Processes

In virtually all processes, the capacities of the various workstations are unbalanced. Goldratt initially developed the production dice game to illustrate the combined effects of dependency and variability on flow processes. Moreover, in *The Goal*, Goldratt (1984) combines insights derived from the basic production dice game and a Boy Scout hike to provide the foundation for understanding the dynamics in unbalanced plants. Models with resources characterized by unbalanced capacity provide the perfect opportunity to introduce concepts such as material release timing, the difference between utilization and activation, categories of capacity (throughput, protective, excess), and the drum-buffer-rope logistical system. For a good discussion of each of these concepts, see Srikanth and Umble (1997).

4.1. Designing Unbalanced Plant Models

The authors recommend modeling an unbalanced plant using an assortment of multi-sided gaming dice. In many hobby/game shops, special gaming dice can be purchased in sets that have 4, 6, 8, 10, 12, and 20 sides. We normally simulate an unbalanced process by setting up plants with six workstations and giving each of the six workers a different type of die to roll.

In *The Goal*, Goldratt (1984) describes three general unbalanced models. The first model occurs when Boy Scouts on a hike essentially arranged themselves from fastest to slowest, with the slowest scout (Herbie) bringing up the rear. The second model is presented when Alex rearranges the Boy Scouts from slowest to fastest. The third model – where Herbie is in the middle of the troop – is described to represent Alex's plant, where the bottleneck resource resides in the middle of the production process. We represent these three different scenarios with models 10 through 12.[3]

Model 10: Decreasing Capacity. In this model, as illustrated in Figure 4a, the six different multi-sided dice are distributed in order of decreasing capacity. That is, the first worker receives a 20-sided die, the second worker receives a 12-sided die, and so on. The average daily capacities for the six workstations are (in order) 10.5 units, 6.5 units, 5.5 units, 4.5 units, 3.5 units, and 2.5 units. As also shown in Figure 4a, starting work-in-process inventory is set at 4 units for each workstation.

Model 11: Increasing Capacity. In this model, as illustrated in Figure 4b, the dice are distributed in order of increasing capacity. That is, the first worker receives a 4-sided die, the second worker receives a 6-sided die, and so on. As also shown in Figure 4b, workstations 2, 4, and 6 start with 3 units of work-in-process inventory. Workstations 3 and 5 start with 2 units of inventory. Since the first workstation is the constraint and will only pass an average of 2.5 units

3. Muth (1977) proved that expected production output is invariant to line reversal under very general conditions. Several other authors, including Gershwin (1987), Conway, Maxwell, McClain, and Thomas (1988), and Dallery and Gershwin (1992) have concluded that, under certain conditions, two tandem queuing systems that are the reverse of one another have the same production rate.

per day through the system, these quantities of inventory are designed to represent a system in equilibrium.

Model 12: Randomly Distributed Capacity. In this model, as illustrated in Figure 4c, the first worker uses a 12-sided die, while workers two through six use a 20-sided, 6-sided, 4-sided, 10-sided, and an 8-sided die, respectively. As also shown in Figure 4c, 4 units of starting work-in-process inventory are allocated to workstations 2, 3, and 4. Since workstations 5 and 6 are located after the constraint resource and will only receive an average of 2.5 units of material per day from the constraint, station 5 starts with 3 units of inventory, while station 6 starts with 2 units.

Figure 4

Figure 4a: Model 10 - Simulating an Unbalanced Plant with Decreasing Capacity

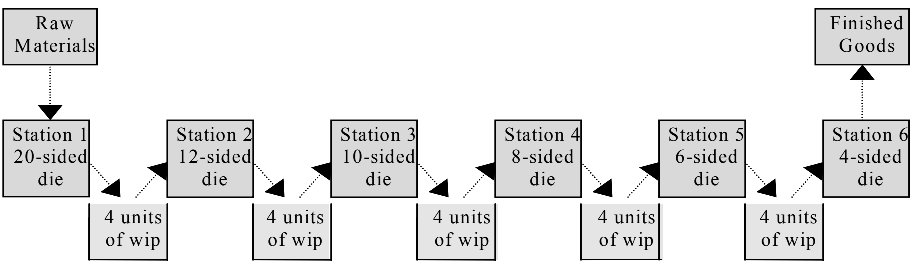

Figure 4b: Model 11 - Simulating an Unbalanced Plant with Increasing Capacity

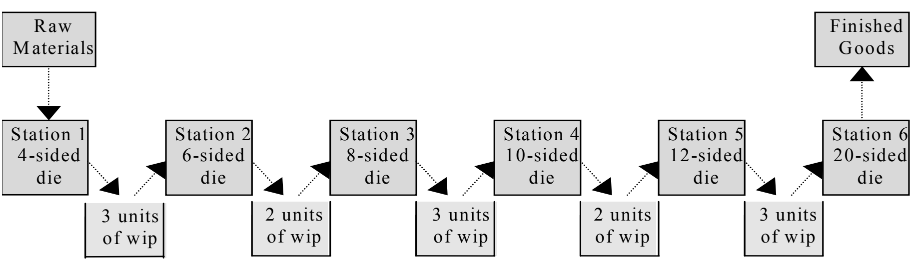

Figure 4c: Model 12 - Simulating a Plant with Randomly Distributed Capacity

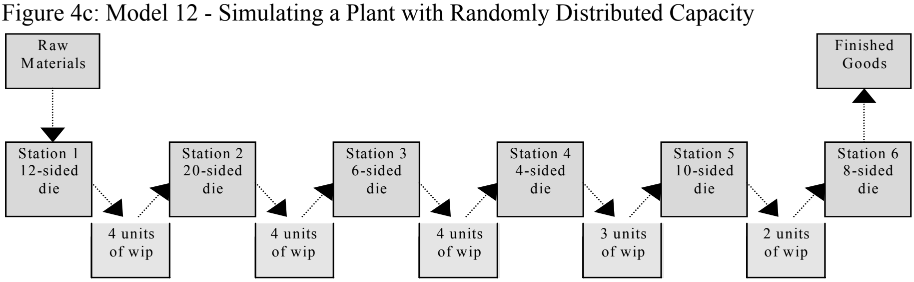

4.2. Pre-Simulation Discussion and Expected Results

Inform the class that they will run their plants for twenty days of operation following the same rules as before. Before performing the simulation, ask each team to predict the following after 20 days of operation: total output for their plant, total ending work-in-process inventory for the entire plant, and which of the workers will have the largest piles of work-in-process inventory. Now you can lead an interesting class discussion where students will make the case for their predictions.

In all cases, the constraint workstation (4-sided die with an average capacity of 2.5 units per day) will limit expected output after 20 days to 50 units. However, when asked, most students will indicate that they do not really expect to get close to 50 units of output. This is because their intuition is heavily biased by their limited experience based on the earlier dependency and variability models for *balanced plants*. In fact, despite the existence of dependency and variability, the output after 20 days will average only slightly less than 50 units, with many of the individual 20 day runs exceeding 50 units. This occurs because all of non-bottleneck workers have some (in some cases significant) amounts of protective capacity.

Model 10 starts with 20 units of work-in-process inventory and should be expected to accumulate an additional 160 units in 20 days. (Station one puts an average of 10.5 units into the system each day and 2.5 units leave station six each day, accumulating system inventory at a rate of 8 units per day.) Analyzing the flow rates, worker 2 will have the greatest amount of inventory because worker 1 processes 10.5 units per day and worker 2 can only process 6.5 units per day. Inventory levels will generally decrease for downstream workstations.

Model 11 starts with 13 units of inventory and should be expected to finish with about the same quantity (since worker 1 puts 2.5 units per day into the system and 2.5 units are expected to be produced each day). The distribution of inventory will not vary greatly, although it should be expected that the level of inventory will slightly decrease the further downstream a worker is located.

Model 12 starts with 17 units of inventory and should accumulate an additional 80 units after 20 days of operation. An examination of worker capacities reveals that inventory will only accumulate consistently at workstations 3 and 4. Station 3 should be expected to accumulate 60 additional units of inventory, while station 4 should be expected to accumulate 20 additional units over the 20 day period.

4.3. Simulation Results

After running the process for 20 days, have each team compile total output, total work-in-process inventory, and individual work-in-process inventory for each station. Compare the actual results to the previously discussed expected results. The students in-class simulation results can then be compared to spreadsheet simulation results that can be generated for each of the three models. The results of the spreadsheet simulations for 1,000 runs of each model are shown in Table 3. The results indicate that the output for models 10 and 12 are statistically significantly higher than for model 11. But the inventory levels for models 10 and 12 are unacceptably high. Now let the class briefly describe how the system should be managed in order to achieve model 10 and 12 levels of output while maintaining model 11 levels of work-in-process inventory, even for, say, a model 12 plant layout. Then let each team run the model 12 configuration using whatever material release and buffering policies they deem appropriate.

Table 3: Results of 1000 Simulation Runs for Models Ten Through Thirteen

	Model 10	Model 11	Model 12	Model 13
Maximum Expected Output	50	50	50	50
Mean Output	48.81	47.82	48.56	48.77
Mean Output Standard Error	(.142)	(.130)	(.136)	(.132)
System Starting WIP	20	13	17	20
Expected System Ending WIP	180	13	97	20
Mean System Ending WIP	181.08	15.17	98.58	21.09
Station 2 Mean Ending WIP	88.86	3.55	8.91	2.30
Station 3 Mean Ending WIP	29.81	3.16	60.81	3.51
Station 4 Mean Ending WIP	23.07	3.00	22.65	9.35
Station 5 Mean Ending WIP	19.71	2.85	2.95	2.82
Station 6 Mean Ending WIP	19.63	2.61	3.26	3.10

4.4. Developing a Fully Synchronized Model

The results from the teams' first run using their own operating rules usually show that output is high but too much inventory is still being released into the system. It is useful to let each team discuss what they did and how well it worked. It is important to make sure that the teams understand that the *material release to the first worker should be carefully controlled*. After a second iteration of running model 12 with their own developed and modified operating rules, most teams will usually arrive at a solution which has two key components. One, the bottleneck station must not be starved. Two, to keep sufficient inventory at the bottleneck without causing excessive amounts of inventory system wide, material release must be synchronized to release materials at the rate that the bottleneck actually consumes material.

The students quickly deduce that a starting buffer of work-in-process inventory large enough to prevent starvation at the bottleneck must be located immediately prior to the bottleneck. The synchronized release of materials to coincide with the bottleneck processing rate is typically implemented in two different ways. Usually, the students first inclination is that since the bottleneck can only process an average of 2.5 units per day, only 2.5 units should be released and made available to station one (by releasing 3 and 2 units on alternate days). However, they soon realize that if the bottleneck works at a faster or slower pace than 2.5 units per day over an extended period of time, then work-in-process inventories at the bottleneck either become too small or needlessly large. They then often devise a feedback mechanism so that whatever number of units the bottleneck consumes becomes the quantity of material released the next day. This mechanism keeps the work-in-process inventory from the material release point to the bottleneck essentially constant.

Table 3 shows the results of a spreadsheet analysis of 1,000 runs of a model labeled as model 13. This model design includes the same layout as model 12, but is managed according to the synchronization guidelines described above with a starting work-in-process inventory of 20 units (fifteen of which are located prior to the bottleneck) with material release averaging 2.5 units per day. The data clearly indicates that a high level of output (statistically the same as models 10 and 12) is achieved with low levels of work-in-process inventory.

5. Conclusion

The production dice game has proven to be a very stimulating and effective classroom learning activity. This paper illustrates how the game can be used in conjunction with spreadsheet simulations to demonstrate a number of key operations management concepts. Student response to these games is excellent and is often very favorably mentioned in student evaluations of the course. The variations described in this paper have simply evolved over a number of years of classroom use. The possible variations of the production dice game is only restricted by the imagination.

References:

Baker, K. R., Powell, S. G. and Pike, D. F. (1990), "Buffered and Unbuffered Assembly Systems with Variable Processing Times", *Journal of Manufacturing and Operations Management* 3: 200-223.

Conway, R., Maxwell, W., McClain, J. O. and Thomas, L. J. (1988), "The Role of Work-in-Process Inventory in Serial Production Lines", *Operations Research* 36(2): 229-241.

Dallery, Y. and Gershwin, S. B. (1992), "Manufacturing Flow Line Systems: A Review of Models and Analytical Results", *Queuing Systems* 23: 3-94.

Garvin, D. A. (1991), "Barriers and Gateways to Learning", from Christensen, C. R., D. A. Garvin, and A. Sweet, *Education for Judgment: The Artistry of Discussion Leadership*, Boston, MA: Harvard Business School Press.

Gershwin, S. B. (1987), "An Efficient Decomposition Method for the Approximate Evaluation of Tandem Queues with Finite Storage Space and Blocking", *Operations Research* 35(2): 291-305.

Goldratt, E. M. and Cox, J. (1984), *The Goal*, Croton-On-Hudson, NY: North River Press.

Johnson, A. C. and Drougas, A. M. (2002), "Using Goldratt's Game to Introduce Simulation in the Introductory Operations Management Course", *INFORMS Transactions on Education* (3)1, http://ite.informs.org/Vol3No1/JohnsonDrougas/

Heineke, J. N. and Meile, L. C. (1995), *Games and Exercises for Operations Management: Hands-On Learning Activities for Basic Concepts and Tools*, Englewood Cliffs, NJ: Prentice Hall.

Muth, E. J. (1977), "The Reversibility Property of Production Lines", *Management Science* 25(2): 152-158.

Ragsdale, C. (2001), "Teaching Management Science with Spreadsheets: From Decision Models to Decision Support", *INFORMS Transactions on Education* 1(2): 68-74, http://ite.informs.org/Vol1No2/Ragsdale/index.php

Srikanth, M. L. and Umble, M. M. (1997), *Synchronous Management: Profit-Based Manufacturing for the 21st Century, Volume 1,* Guilford, CT: Spectrum Publishing.

Umble, M. M. and Srikanth, M. L. (1990), *Synchronous Manufacturing: Principles for World Class Excellence*, Cincinnati, OH: South-Western Publishing Co.

Appendix A: Utilization and Sequencing of Production Dice Game Concepts

The concepts described in this paper can be utilized in undergraduate or graduate core classes or in elective classes that focus on production flows. Clearly the level of the class will determine the degree of theory introduced by the instructor and whether the spreadsheet simulation analysis is given as a problem assignment or if the simulation results are simply presented to the class for reference. The following material describes how the authors utilize the production dice game in their core undergraduate Operations Management class.

We believe that reading *The Goal* gives the students a good foundation for understanding the production dice game concepts. We therefore assign and test our students over *The Goal* and also discuss basic concepts such as throughput, inventory, operating expense, bottlenecks, and constraints before introducing the production dice game.

When introducing the production dice game, we describe the basic setup as illustrated in Figure 1, the logic behind the starting inventory levels, and how material is moved through the system when each worker uses a six-sided die. We calculate the theoretical system output (based strictly on the average capacities) after 20 days of operation. We then ask the class to estimate the actual system output after 20 days and explain the rationale behind their answer. Having read *The Goal*, students intuitively understand that the actual output will fall short of the theoretical output. We use the students' explanations as a segue to define dependency and variability and elicit examples of each in real life processes. We then describe models 2 – 5 which use coins to introduce different degrees of variability into the system. We ask the students to predict the output of each model.

We use six workstations for models 1 – 5 (this sometimes requires having one student run two stations or having one or two "observers") and usually have enough students in the class to run all five models simultaneously. After completing the 20-day simulation we record the results of each team, show the spreadsheet simulation results given in Table 1 and Figure 2, and discuss the results.

Next we introduce models 6 – 9 (which all use six-sided dice but have differing numbers of workstations) and ask the class to predict the output for each model. Similar to the procedure for models 1 – 5, we let the students run the manual simulation and record their results, show the spreadsheet simulation results given in Table 2 and Figure 3, and discuss the results.

Performing the simulations for all of the above models and conducting the corresponding discussion typically requires a full one hour and twenty minute class. In the following class period, unbalanced plant models 10 – 12, as illustrated in Figure 4, are introduced and discussed. The students then run their manual simulations and their results are compared to the spreadsheet simulation results shown in Table 3. The students are then allowed time to run model 12 at least twice using their own policies of material release and buffering strategies with the objective of eventually developing the rules that lead to a fully synchronized model.