5-622-013
MARCH 3, 2022

TEACHING NOTE

Digital Manufacturing at Amgen

Synopsis

Set in January 2020, this case follows Myra Coufal and Chris Garvin, principal scientist and principal engineer, respectively, at biopharmaceutical (biopharma) company Amgen. As members of Amgen's Digital Integration and Predictive Technologies (DIPT) group, Coufal and Garvin help production teams across the company integrate statistical tools into selected work processes.

Amgen manufactures and sells *biologics*, or therapeutics built atop the proteins in living cells. The multi-step biologics manufacturing process is complex, costly, and highly regulated, making many production teams wary of introducing change. Recognizing this reluctance, the DIPT team has prioritized developing statistical models that keep humans in the loop, analyzing data and alerting production teams to potential problems in the manufacturing process. A human decision maker must then decide whether to take action based on the information provided. The DIPT team is beginning to introduce predictive models grounded in machine learning (ML), which are capable of fully automating certain aspects of the manufacturing process.

At the time of the case, the DIPT team has received two requests for support. The first involves working with a team in Amgen's Puerto Rico office to build a standard statistical model for a new commercial product with limited production data. The second involves helping a team in Rhode Island build a predictive ML algorithm to automate one step of the manufacturing process for a top-selling product. Both requests present distinct opportunities and risks. Through analysis and discussion, this case allows students to dissect the complex factors that can impede the success of digital initiatives.

Case Positioning and Purpose

At Harvard Business School (HBS), the Amgen case is used in an elective MBA course called Digital Innovation and Transformation. Cutting across industries, this course aims to equip students to lead and execute digital innovation wherever their careers might take them. The case also is suitable for use in data science and operations management courses, and with undergraduates and executives.

This note was prepared by Professors Shane Greenstein and Kyle R. Myers and Senior Case Researcher Sarah Mehta (Case Research & Writing Group) for the sole purpose of aiding classroom instructors in the use of "Digital Manufacturing at Amgen," HBS No. 621-008. Funding for the development of this note was provided by Harvard Business School and not by the company. It provides analysis and questions that are intended to present alternative approaches to deepening students' comprehension of business issues and energizing classroom discussion. HBS cases are developed solely as the basis for class discussion. Cases are not intended to serve as endorsements, sources of primary data, or illustrations of effective or ineffective management.

At first glance, the Amgen case asks a straightforward question: which opportunity should the DIPT team prioritize? Students can calculate the projected incremental increase in productivity achieved for each opportunity and select the one that yields the highest return on investment (ROI). Deeper analysis, however, reveals several more nuanced considerations that often emerge when companies shift to digital processes (e.g., tackling the statistical limitations of such models, addressing cultural resistance to change, and easing employees' concerns that automation could put them out of work). This case helps students identify, evaluate, and discuss these considerations.

Learning Objectives

The Amgen case aims to:

- Expose students to the unique cost structures of production in the pharmaceutical industry generally, and to the production of biologics more specifically;
- Highlight the difficulties of performing a cost-benefit analysis when predicting the value of investments (namely, in algorithms) that are intended to reduce the uncertainty surrounding rare events that can result in lost batches of production;
- Provide an overview and concrete examples of what "digitization" (or automation) means for manufacturing;
- Explore the spectrum of capabilities of algorithms (e.g., from warning lights that keep humans in the loop to closing the loop with complete automation);
- Evaluate Amgen's challenges and opportunities related to automation for a workforce with different skillsets; and
- Discuss social issues arising from automation.

Supplemental Materials

At HBS, instructors pair this case with a supplemental, quantitative homework assignment (reproduced as **Appendix A**). It elaborates on and provides data related to the central dilemma detailed on case p. 10. Instructors may choose to assign the Amgen case with or without the supplemental homework, depending on their teaching goals and student capabilities.

In addition, instructors may wish to share links to several videos illustrating the Amgen biologics production process. The videos help bring to life the multi-step production process described in the case. They include:

- Maintaining a Clean Environment: https://www.youtube.com/watch?v=MWNP6xgO77U
- Cell Culture: https://www.youtube.com/watch?v=QWhP_CGf6WY
- Purification: https://www.youtube.com/watch?v=98D3x6QeMkg
- Testing: https://www.youtube.com/watch?v=5ok7RumtWzA
- Fill, Finish, and Testing: https://www.youtube.com/watch?v=bzvs_rBR9ug

Suggested Assignment Questions

1. Please address questions 1a, 1b, 2a, and 2b of the homework. Be prepared to discuss your answers at the beginning of class.

2. What factors shape a good experience when developing statistical software for a manufacturing process? What factors shape its subsequent use?

3. What materials do Coufal and Garvin require to develop an ML model? What factors shape its subsequent use by the manufacturing team?

Note: Omit question 1 if the instructor does not plan to assign the homework (**Appendix A**).

Teaching the Case: An 80-Minute Session

At HBS, instructors divide the 80-minute Amgen class session into four main discussion blocks, or pastures, plus a brief introduction and conclusion. These pastures are listed below, along with suggested time allocations.

Introduction	5 minutes
Homework Review	20 minutes
Non-Quantifiable Gains	10 minutes
Development and Implementation of Statistical vs. ML Models	20 minutes
Operational Strategy	20 minutes
Conclusion	5 minutes

This note includes suggestions for teaching the case in either a fully remote or a hybrid classroom. Suggestions rely on standard Zoom features, such as breakout rooms, screen sharing, and polls. **TN Exhibit 1** provides a sample board plan.

Introduction (5 minutes)

Instructors can begin by providing a brief overview of the class session. The discussion will review Amgen's efforts to introduce digital technologies into its manufacturing process, and students will deliberate about which of the two opportunities presented in the case the DIPT team should prioritize.

To begin, ask a student to review the steps involved in the biologics manufacturing process (see case pp. 4-6). To culture and grow the living cells that produce the desired proteins for its products, Amgen mixes raw materials with its proprietary cell lines in small flasks. They then create optimal conditions within those flasks for the cells to grow, divide, and replicate.

As the cells multiply, Amgen transfers them to progressively larger containers, culminating in stainless steel tanks called bioreactors. Once cell multiplication is complete, the company uses a centrifuge to separate the desired protein from excess cell material, purifies the biologic, packages it in syringes, and ships it to wholesalers, who then sell it to pharmacies and hospitals.

The process is delicate, complex, and susceptible to contamination. From start to finish, it takes about two months to complete a batch (case p. 4). While losing a batch is rare, when this does happen it costs Amgen between $1 million and $2 million (case p. 5).

Homework Review (20 minutes)

If the instructor assigned the supplemental homework in advance (refer to **Appendix A**), she can begin the session with a quantitative analysis of the two requests before the DIPT team. If the instructor chose not to assign the homework, she can either pre-board a summary of the answers or skip this pasture altogether, instead devoting the 20 minutes allocated to the homework review to the other pastures. (See **Appendix B** for the full homework answers and **Board 1** in **TN Exhibit 1** for the key points.) *The rest of this discussion assumes that the instructor assigned the homework before class.*

Being by asking a student to provide an overview of the two options. As explained on case p. 10, the DIPT team must choose to either: 1) build a basic, real-time, multivariate, statistical process monitoring (RT-MSPM) model for the mAbX, a fairly new commercial product made in Puerto Rico; or 2) automate one step of the manufacturing process for a top-selling product made in Rhode Island.

For over a decade, Amgen has used RT-MSPM models to help production teams identify anomalies in the manufacturing process (case p. 6). These models analyze data from sensors on the manufacturing equipment and alert teams to any significant deviation from normal parameters (e.g., in batch pressure, temperature, etc.). The team in Puerto Rico has requested an RT-MSPM model to reduce the failure rate of the final five steps of the mAbX production process from 1% to 0.8%. The team in Rhode Island has asked the DIPT team to develop a sensor that would automate a decision about batch quality. This sensor would shave one hour off the total cycle time and would give the DIPT team a valuable opportunity to experiment with models grounded in ML.

Both options present risks. For instance, at the time of the case, Amgen has produced just 10 to 15 small batches of the mAbX. While the company has production data from these small runs, there is no guarantee that these data are representative of production at commercial scale. Coufal and Garvin have limited experience working with the team in Puerto Rico requesting DIPT support and do not know whether this team will be an engaged partner committed to using the RT-MSPM model.

With regard to the second opportunity, the sensor is not perfect. It still produces a small but non-trivial number of false positives and false negatives. A false positive indicates that batch quality is acceptable while in reality, it is not. A false negative indicates that batch quality is poor while truthfully, the quality is fine. These sensor errors could cause Amgen to either waste perfectly acceptable product or move forward with defective product. On the positive side, because the DIPT team has worked with the Rhode Island team several times before, they are confident in the team's commitment to using the sensor once it is developed.

Draw students' attention to question 1a: How much will the plant be willing to spend on research and development to introduce a statistical process that improves the failure rate of each of the last five steps from 1% to 0.8%? Before diving into the quantitative analysis, ask students: *Without using any algebra, just using words, put yourself in Chris's shoes, and explain to your facility manager how this improvement in the failure rate leads to gains to the facility. By what chain of events does it cause lower costs and higher revenue?*

- Reducing the failure rate from 1% to 0.8% lowers the number of failed batches per 100 attempts. Over the course of a year, the number of successful batches will rise and the number of failed batches will fall (or the number of days necessary to produce one successful batch will decline).

- While the material costs rise slightly with more successful batches, a number of these costs are fixed, and are spread out over a larger number of batches. Hence, the average cost per batch declines (slightly).
- In the language of operations, a small reduction in the failure rate increases the output of the facility—what it can produce in a fixed amount of time. In other words, it improves the capacity of the facility over time.

Instructors can now move to the quantitative analysis. Ask students to share how they calculated the plant's annual profits at a 1% failure rate for the final five steps of the mAbX manufacturing process versus a 0.8% failure rate. In short, the annual revenue at a 1% failure rate is $374.83 million, the costs are $362.82 million, and the profit is $12.01 million. When the failure rate drops to 0.8% on each of the last five processes, revenue increases to $379.24 million, and total costs increase slightly to $362.96 million. Thus, at a 0.8% failure rate, total profit equals $16.28 million. This is a net improvement of $4.27 million ($16.28 million - $12.01 million). (For the full analysis, see **Appendix B**.)

Next, the instructor can move to question 1b: It takes one person-year, or 2,000 hours, to develop and refine the model that reduces the failure rate. If PhD statisticians require a minimum of low- to mid-six figure salaries, should Amgen hire a full-time statistician? Take a yes/no vote, and ask students to present their rationale. (If in a hybrid or fully remote setting, use the Zoom polling feature.)

The correct answer is "yes." Over time, the small improvements in the failure rate made possible by the statistician's model add up, yielding more time for the plant to complete successful batches. The improvement in the failure rate generates an increase in the plant's profitability stream in perpetuity, assuming prices stay the same and the firm can sell all its output. The additional gains yield an additional $4.27 million in net profits, and they accrue over time, so the additional statistician is well worth the salary investment. So is another statistician.

Instructors can now move to an analysis of the second request, which comes from a team in Rhode Island that manufactures one of Amgen's top-selling products. This team has asked for DIPT support to install a sensor on a piece of manufacturing equipment that would automate the decision around batch quality—a process that currently takes a person one hour to perform. Ask: *Would you move forward with this request?* Take a yes/no vote, and ask students to present their rationale. (If in a hybrid or fully remote setting, use the Zoom polling feature.)

The main issue with the algorithm in this context is the potential for costly mistakes. These mistakes come in two forms: false positives and false negatives. Because Amgen produces products that ultimately enter patients' bodies, the company is especially concerned about false positives, which mistakenly deem a defective batch acceptable. Senior leaders want to reduce the possibility of false positives as much as possible. To do so, they will accept a large failure rate, even at the expense of some false negatives.

Before moving to the quantitative analysis, ask students to again assume Garvin's perspective. Ask: *Without using any algebra, just using words, put yourself in Chris's shoes and explain to your facility manager how the sensor leads to gains to the facility. By what chain of events does it lower costs and raise revenue?*

The chain of logic is as follows: The algorithm reduces the total amount of time required to process a successful batch. The shorter processing time means that more batches can be attempted in a given time period, BUT the rate of success declines among these larger number of attempts. **On balance,** the financial result depends on the whether the increase in the number of attempts compensates for the success of—and costs incurred due to—the higher failure rate.

Instructors can now ask students to share the results of their quantitative analysis. Question 2 is not as straightforward as question 1. The automated sensor results in a modest loss in profitability (about $130,000 annually). (See **Appendix B** for full analysis.) Therefore, the firm would have to regard the intangible, non-quantifiable benefits generated by the sensor (detailed in the following section) as sufficient justification for the installation. The investment may or may not be worthwhile to the facility.

Instructors can pause here and ask students to notice the continuity of reasoning from the prior question. As with the RT-MSPM model, we can use the basic tools of operations to analyze how the ML algorithm shapes production over the year. Using these tools to predict very low-probability events leads to a reduction in cycle time and increased capacity. We can do this because the sensors try to prevent batch losses, whether or not a human intervenes or a sensor makes an automated decision. As a result, seemingly small progress can generate significant value in this expensive batch setting.

Yet quite a few assumptions go into the calculation, such as assuming that prices and costs stay the same, and making an educated guess about the failure rate of the algorithm. We also do not know how long it will take to develop the algorithm and the challenges that might emerge as we install it. These uncertainties lower the confidence level in our predictions. If these insights do not emerge organically in the discussion, instructors can share them explicitly.

Non-Quantifiable Gains (10 minutes)

Instructors can transition to the next discussion pasture by broadening the discussion around option 2. Now that the quantitative analysis is complete, ask students to consider the non-quantifiable benefits of automation, as laid out in question 2b. (Note: If the instructor did not assign the homework before class, she can still pose this question.) Ask: *What tactical advantages does the facility gain from using an ML algorithm to perform the quality test? How would you explain these advantages to the facility manager?* Record student responses on **Board 2** (see **TN Exhibit 1**). Responses typically include:

- Automating the quality test reduces cycle time, in turn increasing facility capacity. Over time, that excess capacity could:
  - Delay or prevent the need for capital outlays to increase plant capacity or build a new plant;
  - Avoid the cost of eventually needing to outsource to a contract manufacturing organization for extra capacity;
  - Increase schedule flexibility for plant maintenance, upgrades, etc.;
  - Offer the plant manager the option of performing additional checks between batches, which could result in fewer errors and a higher batch success rate;
  - Help to accommodate unexpected demand surges; and
  - Enable the plant manager to dedicate manufacturing capacity to another product, which could help speed commercialization if that capacity is allocated to new products.
- The time and labor that used to go into quality testing can be reallocated to other tasks.
- Learning how to automate effectively could help the DIPT team perform later experiments. The lessons learned could prove useful for installing sensors at an earlier point in the manufacturing process, thus saving additional time.
- Detecting production input material errors sooner in the process leads to earlier conversations with suppliers and reduces future quality control issues.

Conclude this discussion by recapping options 1 and 2. The analysis of option 1 showed that a reduction in the failure rate leads to an increase in the number of successful batches per year, which directly affects plant costs and revenues. The analysis of option 2 showed that automating a sensor shaves time from the overall production process and increases total production over the course of the year. However, it also increases the batch failure rate, so it is not an overall winner in every circumstance. These calculations begin to hint at the complexity of automation.

Development and Implementation of Statistical vs. ML Models (20 minutes)

Next, instructors can ask students to consider how challenging it might be for the DIPT team to develop and implement the RT-MSPM model and the ML algorithm. Begin with the RT-MSPM model. Ask: *What factors make it easier or more challenging to develop and introduce the new sensor into the mAbX manufacturing process in Puerto Rico?* Responses may include:

- Amgen has been using RT-MSPM models since the early 2010s. Staff are familiar and comfortable with them. As of 2019, the DIPT team has deployed 800 RT-MSPM models covering 1,000 operations across 14 manufacturing plants (case p. 7). This bodes well both for the acceptance of the new model and the DIPT team's ability to build the model.
- As noted in the case (p. 7), buy-in among local production teams is critical for sustained use of RT-MSPM models. The DIPT team has considerable experience working with teams in the Puerto Rico office. They feel confident that the plant's leadership will support and encourage adoption of the new RT-MSPM model. In addition, before moving to Massachusetts in 2015, Coufal spent a decade working in Puerto Rico, so she has a number of connections and relationships in the office on which to build (case p. 10). The Puerto Rico office was also home to the DIPT's first attempt to automate a manufacturing decision (case p. 9), indicating a cultural willingness to experiment and take risks. However, Coufal and Garvin have never worked directly with the mAbX team, which introduces some level of adoption/engagement risk.
- The limited mAbX production data poses statistical risk. As the case notes (p. 10), Amgen has produced 10 to 15 small batches of the product and five batches at commercial scale. RT-MSPM models are based on historical production data that they use to identify normal pressure, temperature, and pH parameters for products (case p. 6). The less production data available, the less confident the DIPT team will be in the accuracy and precision of the RT-MSPM model.
- A company-wide challenge is "alarm fatigue." As the case notes (p. 7), a number of RT-MSPM alarms are not indicative of an actual problem, which makes people less sensitive to them over time. Alarm fatigue could reduce long-term effectiveness of the mAbX RT-MSPM model.
- The 2,000 staff members working in the Puerto Rico office have varying levels of education and experience. Just over half (56%) hold a bachelor's degree, 18% hold an advanced degree, and the remaining 26% have no education beyond high school (case p. 4). The DIPT team will need to convey the RT-MSPM's value proposition in terms that resonate with these different employee populations.

Next, move to option 2. Ask: *What factors will make it easier or more challenging to develop and introduce the ML algorithm into the Rhode Island production facility?* Responses may include:

- The DIPT team has worked extensively with this production team in Rhode Island. Thus, Coufal and Garvin feel confident about the team's sustained commitment to using the sensor.
- Amgen's production teams are reluctant to remove humans from the loop. Since losing a batch of product is very costly, plant managers often feel more comfortable controlling the production process themselves. While human judgment may be equally or even more fallible than the

sensor, manufacturing staff often do not perceive this to be the case. The reluctance to automate tasks could hamper the DIPT team's efforts.

- The team in Rhode Island is looking to automate the quality test for a "top-selling product." While the case does not name the product in question, astute students deduce that it must be either Enbrel®, Neulasta®, or Prolia®. As noted in case **Exhibit 2**, these three products generate 50% of Amgen's sales. Given the importance of these three products, the stakes are high. If the automated test causes Amgen to lose a substantially higher number of batches, the issue would affect a major driver of revenue.

- The case does not specify whether Amgen will need to amend its regulatory filing with the U.S. Food and Drug Administration (FDA) if it automates the quality test, but this seems like a safe assumption. As noted in the case (p. 3), regulatory filings include exhaustive detail about how products are to be developed; automating the batch quality test would likely require Amgen to go back to the FDA. This entails substantial time and cost.

- Finally, it is reasonable to assume that at least a few manufacturing staff will push back on the automated sensor for fear of it being the first step to replacing their jobs. The DIPT team and its counterparts in Rhode Island must acknowledge employees' concerns and take care to thoroughly convey the benefits of digital technologies.

At this point, instructors can broaden the discussion to other contexts. Ask: *Why do we call manufacturing in biotech a "batch" production method? What does "batch" mean in this context?* If in a remote or hybrid classroom setting, use the share screen function to show case **Exhibit 5**, which depicts Amgen's biologics manufacturing and packaging process. The biologics manufacturing process is batch production because it consists of a number of sequential steps, all meant to produce an identical set of products at one time. This is similar to batch production in other industries, which means that we can generalize lessons learned from Amgen to other batch settings.

Given this, ask students to imagine that their manager has tasked them with overseeing the deployment of statistical models, like RT-MSPM, in a manufacturing facility in their future career (not necessarily in biopharma). Ask: *What factors could make such a program go well? What would you do to set yourself up for success at three stages: 1) when developing the models; 2) when initially installing the sensors; and 3) when extending the use of sensors?* Record responses on **Board 3** (see **TN Exhibit 1**).

- Development of models: Bottom-up model development—the impetus should come from manufacturing staff in response to specific business problems. As Steve Fuller and his team did at Amgen (case p. 7), workers from the floor should aid in model development to increase buy-in and familiarity with the model. Digital integration teams should anticipate and mitigate reluctance to adopt models as much as they can.

- Initial installation of sensors: Again, bottom-up process—facility management and staff should support and advocate for the use of the sensors and be willing to experiment and take risks. Sensor installation should not take too much time. Lastly, to demonstrate viability, initial sensors should be added to a process or piece of equipment that poses relatively few risks.

- Extending the use of sensors: ideally, workers on the facility floor should view the sensors as a tool to free their time for other, more valuable work. Facility employees should have positive early experiences with the sensors; this helps to create trust and build excitement for their extended use. Workers at the plant should be willing and able to learn new skills.

Next, instructors can ask students to imagine that their manager in this same hypothetical future manufacturing facility has tasked them with overseeing the deployment of an ML algorithm. The facility already has some experience with sensors. What additional factors come into place now that could make a program for automation go well? Ask: *What additional steps would you take to set yourself up for success at three stages: 1) when developing the algorithm; 2) when initially installing the sensor; and 3) when extending its use?* Record responses on **Board 4** (see **TN Exhibit 1**).

- Development of the algorithm: In addition to all the conditions noted in the preceding question, closed-loop algorithms (i.e., no human involved) should be relatively simple. While the processes that go into ML models are often a black box (case p. 3), it is important to be able to clearly explain the algorithm's purpose. Facility management and digital integration teams should carefully consider the costs of false negatives and false positives to understand the implications of each. Doing so requires a deep familiarity with the production environment and an ability to iterate with a committed production team for a long period. Optimally, the data used to develop one ML algorithm should feed into training sets for other algorithms.
- Initial installation of sensors: In addition to all the conditions noted in the preceding question, facility management and the digital integration team should test the algorithm under a variety of settings to understand whether it accurately captures false positives and false negatives. As the sensors are being installed, the digital integration team should emphasize that the algorithm can eliminate operator error. The extent to which employees care about reducing this is unclear, however.
- Extending use of sensors: In addition to all the conditions noted in the preceding question, facility management should clearly communicate to manufacturing floor workers that their job duties and employment status are not at risk. Optimally, the facility should learn from experimentation as it extends use of the sensors. Lastly, implementing the sensors successfully requires considerable trust built between the statisticians and the production team.

Instructors can conclude this discussion pasture by pausing to reflect on the issues surfaced. Note that all of these considerations were spurred by the seemingly simple decision to add a sensor to a piece of manufacturing equipment. The key insight is that the addition of a statistical process involves more than just calculating a few statistics, installing the sensor, and waiting to take advantage of the benefits. It involves a number of iterative steps in the sensor's development and implementation. In some manufacturing facilities, implementation will be easier than in others. Indeed, the payoff from model development and installation depends on the interaction with the setting. The importance of engaged, supportive leaders cannot be overstated.

Operational Strategy (20 minutes)

Examine this insight further by shifting to the perspective of a factory worker. Tell students: *Put yourself in the shoes of a line worker who has been at Amgen for 10 years. Chris's team is now coming in for the fifth time to automate something that you used to control. How would you respond?* It is likely that the line worker would feel frustrated by the constant change and perceived redundancy of his or her tasks. As more and more sensors come into play, the line worker becomes increasingly concerned about his or her job security.

With this perspective in mind, divide students into groups of four for a five-minute discussion. (If in a remote or hybrid setting, use the Zoom breakout room function here.) Give the groups the following discussion prompt: *What challenges will shape the success of an operational strategy for implementing digital technologies across all of Amgen?* When students return from breakout rooms, ask a few to share their insights. Responses typically include:

- Companies need a plan of action and expertise in building and implementing digital technologies, and leadership MUST be committed to supporting the effort. For example, this type of program requires commitment to employees who are skilled at statistical modeling, investments that may lose money at first, and programs to share lessons learned across installations.

- A centralized team should be responsible for calculating the ROI in order to make tradeoffs between potential opportunities.
- Companies should identify the parties responsible for cataloguing accumulated lessons, carrying lessons across facilities, and addressing periodic and ongoing issues.
- Companies must account for the human element. Framing ML as an enhancement to workers is important to ensure acceptance by all staff.

If the discussion stalls, instructors can use the below questions to dig deeper:

- *Push Question 1*: How does the scope of potential applications shape operational strategy? What percent of the firm's parts (e.g., divisions, plants, production lines) are involved in any degree of automation?
- *Push Question 2*: How does the scale of the program shape operational strategy? What does a program aspire to accomplish within a given part of the firm, and how much of those employees' current responsibilities are being automated?

Students should recognize that while sensors are widespread, there is still plenty of room for even more sensors in some plants. Amgen has just begun to experiment with automation; it is nowhere near the possibility of replacing any worker. If done well, the push to automate will begin to relieve alarm fatigue and potentially free up a modest amount of time for them to do other tasks.

Often, students conclude that, as Amgen scales its digital manufacturing efforts, the company might benefit from hiring someone to accumulate lessons across production teams and frame the tradeoffs for those facilities interested in digital manufacturing. At this point, instructors can reveal that in April 2020, case co-protagonist Myra Coufal assumed the role of Principal for Digital Operations.[1] In this new role, she does just that.

Conclusion (5 minutes)

By the end of the session, students should understand that simply performing a cost-benefit or ROI calculation is insufficient for making decisions regarding digital innovation. Rather, teams like the DIPT must work to address the statistical limitations of models, cultivate a company culture and mindset amenable to experimentation, and actively demonstrate the benefit digital technologies can offer to employees. Absent this effort, digital innovation efforts are often short-lived.

Conclude the class session by asking a student to summarize lessons learned. On the surface, this has been a conversation about when and where organizations choose to implement ML in manufacturing. Discussion begins with a quantitative analysis of changes to Amgen's costs, revenues, and profits stemming from increased capacity and reduced waste. But deeper analysis reveals a number of additional, less obvious considerations.

To succeed, Amgen must elicit the right actions/choices among plant managers, build maximum staff buy-in, address automation concerns, and generate and capture data and insights. These considerations must be integrated into a company's broader operations strategy.

[1] Myra Coufal, LinkedIn profile, https://www.linkedin.com/in/myra-coufal-9669a5175/, accessed January 2021.

TN Exhibit 1 Sample Board Plan

Board 1: ***Homework Answers***

Question 1a:
- When the failure rate drops from 1% to 0.8%, annual net income increases by $4.27 million.

Question 1b:
- Hire the statistician: the value generated far exceeds their salary.

Question 2a:
- The sensor yields a modest loss in profitability. It might still be worth the expense of installing it, however, if the statistician was too pessimistic in their estimation of the failure rate.

Question 2b:
- See **Board 2**.

Board 2: ***Non-quantifiable gains of automation***

- The time and labor that used to go into quality testing can be reallocated to other tasks.
- The incremental capacity could:
  - be allocated to manufacturing another product;
  - avoid outsourcing manufacturing;
  - delay/avoid capital costs;
  - allow faster product commercialization; and
  - cover unexpected demand surges.
- Learning how to automate effectively could help with later experiments.
- There are gains to saving on wasted material.
- Potentially useful for installing sensors at earlier points in the process, and better measurement detects material errors sooner.

Board 3: ***Optimal conditions for introducing statistical models such as the RT-MSPM***

<u>Development of Models:</u>
- Bottom up: factory staff seek the opportunity out.
- Floor workers give their time to statisticians to help develop models for sensors.
- The statistical problems resemble statistical problems the statisticians have seen in previous programs in other facilities, so development of statistical model is easy.
- Actively anticipate that some workers will be resistant to the models.

<u>Initial Installation of Sensors:</u>
- Bottom up: factory staff seek the opportunity out.
- Facility management is enthusiastic and actively champions the program.
- Does not need too much pre-launch testing.
- Have equipment on hand that can demonstrate viability early at low risk.

<u>Extending Use of Sensors:</u>
- Floor workers perceive sensors as something to free their time for other, more valuable work.
- Installing a sensor on one tank illustrates how to approach installation on other tanks.
- Early experience goes well, earning trust of facility employees, so later activities get their cooperation.
- Broad willingness among floor workers, managers, and modelers, to learn new skills, or adapt existing skills to the new situation.

Board 4: ***Optimal conditions for introducing machine learning algorithms***

<u>Development of Algorithm:</u>
- Again, bottom up.
- Again, floor workers help develop models.
- Again, some staff will resist.
- NEW: The models for the closed loop cannot be too complicated.
- NEW: Must carefully consider the costs of false negatives and false positives.
- NEW: May need to iterate for a longer period of time.
- NEW: Data in one project may be useful for training sets in another.

<u>Initial Installation of Sensors:</u>
- Again, bottom up.
- Again, engaged management and local advocates.
- Again, does not take too much time.
- Again, willingness to experiment and take risks.
- NEW: Need long time to test the model under a variety of settings to understand whether the ML accurately captures false pos and negs.
- NEW: Eliminates operator error. How focused are employees on reducing?

<u>Extending Use of Sensors:</u>
- Again, floor workers perceive sensors as valuable.
- Again, willingness to learn.
- Again, early experiences must be positive.
- NEW: Floor workers do not perceive their job duties or employment status as at risk.
- NEW: Due to the novelty of the activities, experimentation should have additional value for later experiments.
- NEW: Risky for everyone due to unknown unknowns. It is not possible without trust.

Source: Casewriters.

Appendix A: Optional Homework Assignment

Option 1: Building an RT-MSPM Model for the mAbX product

Goal: Reduce the batch failure rate. (Note: Reducing the batch failure rate is only one possible outcome; the model would also likely detect additional anomalies throughout the manufacturing process.)

Setting: A team in Puerto Rico has requested a statistical model to reduce the failure rate of a batch of mAbX. Each batch yields product worth $3.25 million in sales.

A plant runs at full capacity if it operates 365.25 days per year (the extra 0.25 comes from leap years). Assume that the plant can sell everything it produces.

Production for each batch consists of 10 steps and in total takes three days to create. The last five steps are delicate, and each has a 1% failure rate: 1 in 100 batches fail at one of these points in the process. Failures in one step are statistically independent of every other step, so the failure rate for a batch is 5%. It is not possible to tell whether the entire batch is lost until the end. A quality assurance test tells the plant manager whether the batch quality is acceptable or unacceptable.

Each of the 10 steps takes seven hours. The quality assurance test must be performed in a lab and takes one hour. After a successful test, all equipment is cleaned, which takes one hour. An entire successful cycle takes 72 hours, after which the cycle can start again. After a failed test, all equipment is cleaned for 5 hours. An entire unsuccessful cycle takes 76 hours. After cleaning, the cycle starts again.

Some costs are incurred no matter what the plant produces (i.e., plant labor, electricity, maintenance, plant security, debt for the structure, etc.); we call these overhead costs. The overhead costs are $120 million per year. Other costs depend on production; these are called variable costs. Each step of the manufacturing process incurs $200,000 of material costs, so a batch costs $2 million ($200,000 per step*10 steps).

Question 1a: How much will the plant be willing to spend on R&D to introduce a statistical process that improves the failure rate of each of the last five steps from 1% to 0.8%? Break it down to several distinct questions. How much does capacity utilization change? How much cost would be saved on wasted material? How much additional profit would the plant generate?

Question 1b: It takes one person-year, or 2,000 hours, to develop and perfect the model that reduces the failure rate. If PhD statisticians require a minimum of low-to-middle six-figure salaries, is it ever worthwhile to hire a full-time employee? When is it not? What if it took 4,000 hours instead, or two full-time employees? Is it still worth hiring both employees?

Option 2: Automating One Step in the Manufacturing Process for a Top-Selling Product

Goal: Automate a batch quality test.

Setting: Assume that the manufacturing process for the top-selling drug in question is similar to the process outlined in Option 1. Once again, start from a setting with 10 processes, where the last five each display a failure rate of 1% (for a total failure rate of 5%). Once again, the entire success cycle takes 70 hours, plus one hour of testing for quality and one hour of cleaning, and the unsuccessful cycle takes 70 hours, plus one hour of testing and five hours of cleaning.

At present, it is somebody's job to do these tests for quality. The statisticians propose installing an additional sensor on already extensive array of sensors on the manufacturing equipment. The sensor automates the decision around quality. It costs $500,000 to purchase and install.

Adding this new sensor would eliminate the one-hour wait for the lab results. It would reduce the time for each successful batch from 72 to 71 hours. The time for an unsuccessful batch also declines from 76 to 75 hours.

The sensor needs a model to perform the equivalent of a test for quality. That model uses machine learning. It provides a decision—green/red for success/fail. The model attempts to replicate what the person would have decided after viewing the lab results. Like all machine learning models, however, these automated tests struggle to precisely replicate human judgment. They produce a tiny but non-trivial rate of false positives and a tiny but non-trivial rate of false negatives. False positives give a green light even though the batch spoiled, and false negatives give a red light even though the batch truthfully was unspoiled.

Statisticians can set parameters in the decision rules that determine the green/red to increase or decrease the probability of false negatives and false positives. Because the underlying biochemistry is so complex, there is a trade-off between false negatives and false positives, whereby the parameters that lower the probability of one increase the probability of the other, and vice versa.

The firm's senior management believe that their reputation with their customers crucially depends on never delivering ineffective output. That belief aligns with government regulations in many countries, which heavily penalize firms for shipping sub-par products. Therefore, the plant's managers instruct the statisticians to reduce the probability of false positives to its lowest possible point. Due to this change, the statisticians estimate that installing the automated sensor will increase the failure rate to "more or less 5.5%." They cannot be sure until they operate for a while.

Question 2a: What are the quantifiable gains from installing this sensor? Based on the numbers, is it worth the expense?

Question 2b: Improvements in cycle time generate a range of gains to the firm that cannot be quantified. List gains from automation for the operating strategy of the firm, whether quantifiable or not, and list as many as possible.

Source: Casewriters, based on input from the company.

Appendix B: Homework Answers

Question 1a: How much will the plant be willing to spend on R&D to introduce a statistical process that improves the failure rate of each of the last five steps from 1% to 0.8%?

Answer

Revenue at 1% failure rate: On average, we can expect five failures for every 100 attempts, so 95 successful attempts will take 72 hours each and five failed attempts will take 76 hours each. That totals to 95*72 + 5*76 = 6,840 + 380 = 7,220 hours for 95 successes, which implies an average of 76 hours, or 3.167 days, per successful batch. That extra 0.167 day per successful batch accounts for the average time spent on failures. That implies 365.25/3.167 = 115.33 successful batches per year. Thus, in three out of every four years, roughly, the plant will get 115 successful batches, and in one out of every four years, roughly, it will get 116. Notice that also implies that 115.33*72 hours = 8,303.76 hours, or 345.99 days, are spent on successful batches in a year, while 19.26 days are spent on unsuccessful batches. That implies an average of 6.08 unsuccessful batches per year. In total, the plant attempts 115.33 + 6.08 = 121.41 batches per year. The plant makes revenue only on successful batches, so expected revenue is $3.25 million*115.33 = $374.83 million.

Costs at 1% failure rate: What are the total expected costs? Overhead costs are incurred on each batch, whether it succeeds or fails. Total overhead for the year = $120 million. The costs of materials are incurred whether the batch succeeds or fails. Costs of materials are $2 million*121.41 attempted batches= $242.82 million. In total, the plant can expect annual costs of $120 million + $242.82 million = $362.82 million.

Profit at 1% failure rate: Profitability is total revenue less total costs, or $374.83 million - $362.82 million = $12.01 million.

Improvement: What happens when the failure rate drops to 0.08 on each of the last five steps? In 100 attempts, the lower failure rate leads to 5*0.8 = 4 failures on average, increasing the success rate to 96 batches out of 100. Those take 72*96 + 76*4 = 6,912 + 304 = 7,216 hours for 96 successes. That implies 3.13 days per successful batch (7,216 hours/96 successes/24 hours), an improvement over the previous time of 3.167 days. Annually, this seemingly small improvement means that the plant can produce 365.25/3.13, or 116.69 successful batches per year at a failure rate of 0.8%, an increase over the 115.33 successful batches per year when the failure rate is 1%. This also implies that 350.07 days per year are spent creating successful batches (116.69 batches*72 hours per batch/24 hours), while 15.18 days are spent creating 4.79 unsuccessful batches (a decline from 6.08). At a 0.8% failure rate, 121.48 total batches are produced per year, an increase over the 121.41 produced at a failure rate of 1%. The increase in the total number of batches comes from the smaller cumulative amount of time needed to clean due to the smaller number of unsuccessful batches (four at the 0.8% failure rate vs. five at the 1% rate).

Revenue and Cost with Improvement: Revenue now rises to $3.25 million*116.69 batches = $379.24 million. That is an improvement of $4.41 million ($379.24 million at a 0.8% failure rate - $374.83 million at a 1% failure rate). Annual overhead remains unchanged ($120 million). Material costs increase to $2 million*121.48 batches = $242.96 million (a slight increase from $242.82 million). Total costs become $120 million (overhead) + $242.96 million (materials costs) = $362.96 million, a total increase of $140,000 from the $362.82 million in total costs at a 1% failure rate. On balance, that is a small cost increase compared to the increase in revenue, leaving a net improvement of $4.27 million. The gain comes from increased revenue due to improved plant capacity, which comes from more efficient use of its time.

Question 1b: Is it worth hiring the statistician? How about two statisticians?

Answer:

You may have thought a small change in failure rate would not be valuable. Over the course of a year, however, the small change adds up, yielding more time for successful batches. The reduced failure rate yields an increase in the plant's profitability stream in perpetuity, assuming that prices stay the same and the plant can sell all its output. In principle, that is quite a bit of new improvement, depending on how you discount it. Will one person cost that much for one person-year? No. The statistician is paid in the low- to mid-six figures. How about two people? Again, that is far lower than the $4.27 million increase in net annual profits. In conclusion, yes, Amgen should hire the statistician.

Question 2a: What are the quantifiable gains from installing this sensor? Based on the numbers, is it worth the expense?

Answer:

At face value, the increased failure rate of 5.5% is problematic, but the automated testing shaves one hour off cycle time, which could be a net positive. Whereas successful batches previously took 72 hours to create, they need 71 hours using the automated sensor; unsuccessful batches take 75 hours instead of 76. This means the total hours required in 100 cases is 71*94.5 + 75*5.5 = 6,709.5 + 412.5 = 7,122. That means it takes 3.14 days per successful batch (7,122 hours/94.5 successes*24 hours), meaning that 365.25/3.14 = 116.32 successful batches are created per year. That implies 116.32*71 hours = 8,258.72 hours, or 344.11 days, are spent to create successful batches, while 21.14 days are spent creating 6.76 unsuccessful batches. Notice the total number of attempts increases to 116.32+6.76 = 123.08.

Revenue and Cost Changes: The plant can expect to generate $3.25 million*116.32 successful batches = $378.04 million in annual revenue. That is a $3.21 million increase over $374.83 million (total revenue from question 1a above). What are the total expected costs? Overhead costs are still $120 million per year, but the total number of batches has increased, so material costs have increased with it. Those costs are $2 million*123.08 = $246.16 million, an increase of $3.34 million over $242.82 million (total costs from question 1a above). Thus, while annual net profits at a 1% failure rate were $12.01 million ($374.83 million - $242.82 million - $120 million; see question 1a above), adding the sensor reduces profits to $11.88 million ($378.04 million - $246.16 million - $120 million), a $130,000 decline.

Analysis: When the automated sensor is installed, two things happen: revenue increases due to the additional plant capacity, and costs also increase due to the additional materials needed to produce extra batches. On balance, the sensor reduces annual profits by $130,000. Despite this, it might still be worth the expense of installing the sensor because there are several unknowns at play. The statisticians might have been too pessimistic in their estimation of the failure rate. If this is the case and the failure rate is actually lower, automation could be highly profitable. But if the failure rate is actually worse than the estimation, the plant could lose even more money than projected. The statisticians should watch the early experience carefully to make certain the losses are not too high.

Question 2b: Improvements in cycle time generate a range of gains to the firm that cannot be quantified. List as many gains as possible for the firm's operating strategy stemming from automation.

Answer:

See **Board 2** in **TN Exhibit 1** for a detailed list of non-quantifiable gains.

Source: Casewriters.