



Class Notes: BPM Research and Education

Michael zur Muehlen

Director, Center of Excellence in Business Process Innovation
Howe School of Technology Management
Stevens Institute of Technology

Michael.zurMuehlen@stevens.edu
<http://howe.stevens.edu/BPM>

Evidence-based Research: What we know, what we don't, and why we should find out

Would you let your doctor write a prescription because he has a hunch that it might be the right one? Malcolm Gladwell's bestseller *Blink* has advocated that intuition is a great indicator when employed by a skilled expert; it allows us to focus on the core features of a situation with limited cognitive effort. So, would you accept the doctor's hunch? I guess not. We prefer to have our doctors base their decisions on the results of successful clinical trials. These trials should have applied the treatment in question to patients in a situation comparable to ours. The underlying principle is called *evidence-based medicine*. While it does not apply in all areas of medicine, making decisions based on data provided by sound research is the gold standard in many areas.

The Agency for Healthcare Research Quality has released a classification of data that rates data according to the way it was obtained. Evidence at Level I has been obtained using the most rigorous scientific methods, while evidence at Level III is typically based on experience and conjecture.

- Level I: Evidence obtained from at least one properly designed randomized controlled trial.
- Level II-1: Evidence obtained from well-designed controlled trials without randomization.
- Level II-2: Evidence obtained from well-designed cohort or case-control analytic studies, preferably from more than one center or research group.
- Level II-3: Evidence obtained from multiple time series with or without the intervention. Dramatic results in uncontrolled trials might also be regarded as this type of evidence.
- Level III: Opinions of respected authorities, based on clinical experience, descriptive studies, or reports of expert committees.

This evidence-based approach forces us to revisit what we know, what we think we know, and where we have to rely on guesswork.

Business Process Management aims to improve the operation of organizations. It cures aching joints caused by the misalignment of technology and business operations, fixes clogged process arteries by separating decision making from transaction processingⁱ and establishes a healthy exercise regimen through the continual review and improvement of process performance. So how do we make decisions about the design and implementation of our Business Processes? And how do we know we are doing the right thing?

Don't like your Process Modeling Language? Invent a new one.

We know surprisingly little about how people model business processes. Flowcharts have been around since the 1930s, IDEF since the late 1970s, UML since the early 1990s, and BPMN since 2006 (and unofficially since 2004). Yet, we know from experience that the process diagrams of

novice modelers differ from those of experienced consultants. This is true not just for process modeling, but for other areas of conceptual modeling as well.ⁱⁱ A lot of energy is devoted to the development of new modeling languages and techniques. In 1998 the Process Specification Language project at NIST conducted a review of then current process modeling languages.ⁱⁱⁱ Twenty-six representations were identified and analyzed. Since then, BPMN was ratified, and numerous other process-modeling languages have been proposed. A typical justification for the development of a new language is that existing languages are seen as not expressive enough. “I cannot say what I want to say using language X” or “I need to model X, but the existing languages do not provide a construct for this” are the main drivers behind this. But what is the level of evidence we use to come to this conclusion?

The tendency to invent new process modeling languages is akin to learning a new language whenever you cannot precisely express what you want to say in the language(s) you speak. Can't say “Ohrwurm”^{iv} in English? English must be bad. Let's use Esperanto. The result? A Tower of Babel where few people communicate using the same language. And a thriving market for translators and multi-linguists.

Lingua Franca: BPMN

There is an increasing level of evidence (Level III, to be exact) that BPMN is becoming the lingua franca of the process modeling community. Compared with other process modeling languages, BPMN has a large vocabulary (52 symbols in version 1.1). More than 100 people participated in more than 140 meetings (face-to-face and virtual) to hash out the details of the BPMN specification. If we compare BPMN with a natural language, three questions come up:

1. How can BPMN be used (i.e., what's possible in theory)?
2. How should BPMN be used (i.e., what's recommended for practice)?
3. How is BPMN being used (i.e., what do people actually do with it)?

Finding answers to these questions is significant. The BPMN specification provides a list of constructs (the vocabulary) and how they may be connected (the grammar). This is equivalent to an English dictionary that contains the meaning of words and rules for how to build sentences. It does not give an indication of the content coverage provided by the language; that's hard to do for a natural language, but easier for a modeling language with a restricted vocabulary.

A dictionary does not provide a guide for what constitutes an elegant sentence or a colloquial one (although it might indicate typical word usage). In English that's where style guides such as Zinsser's *On Writing Well* or Strunk & White's *Elements of Style* come in. In BPMN no such universal advice exists. We have to rely of design guidelines that are specific to organizations or have been developed by trainers and consultants.

Finally, to improve BPMN literacy, we need to understand what the current use of the language looks like, what mistakes modelers typically make, and which of the recommended language areas are currently underutilized. This is akin to documenting the as-is state of business processes. While you know that the processes will change, knowing what they look like now enables you to focus on the areas that will change, and plan for the direction in which they will change.

What We Know

Possibilities of BPMN

The first question can be answered by looking at the BPMN specification and by coming up with metrics for the expressiveness of the notation. A lot of academic work has been published in this area. BPMN has been compared, using workflow patterns^v, evaluated directly against other process modeling languages^{vi}, and its constructs (and those of other languages) have been mapped against benchmark ontologies^{vii}. We know that BPMN fares fairly well when compared to

other process modeling languages. It provides a rich vocabulary, mechanisms to represent asynchronous communication among processes, and has been designed with a mapping to BPEL in mind. Its richness is also one of the weak points: There are multiple ways in which a modeler can express the same situation. A decision in a process, for example, can be represented implicit, using conditional flow following an activity, or explicit, using a gateway shape. Exclusive branching can be expressed using a blank diamond, or a diamond marked with an X. The impact of this symbol overlap can be interpreted differently. Some authors indicated they would prefer a leaner language where there is just one “official” way to express a situation. Others have interviewed users who indicated that this situation is seen as a sign of language flexibility, and, thus, a good thing. Our evidence in this area is at Level II-3 to II-2.

Proper Use of BPMN

The second question is based on the normative decision what constitutes “good” process modeling practices. This is the domain of BPMN developers, modeling practitioners, and trainers. The BPMN specification contains some (but not much) advice on how the individual symbols should be used. BPMN trainers such as Bruce Silver and Derek Miers have put a lot of thought into teaching good BPMN modeling practices, and Bruce has published much of his excellent work on his blog.^{viii}

On the academic side, some work has been published on generic Guidelines of Modeling^{ix}, but these are not specific to BPMN. The Guidelines of Modeling state that models have to be correct, relevant to the problem domain, and that their creation has to be economically feasible. They should be clearly laid out, follow a systematic design, and be comparable to other models. These generic guidelines have to be translated into actionable advice for BPMN modelers. For example: A pool representing the customer should always be at the top of the diagram. Activities in one diagram should have a comparable level of granularity. Models should be arranged within a larger framework that organizes them and places them in context. Modelers should choose one representation for gateways and stick with it.

To evaluate the practicality and quality of these guidelines, we would need to establish how easy the resulting models are to understand and whether they represent their content with a minimum amount of complexity. This is difficult to measure since complexity relates to the number and variety of symbols in a diagram and to the labels of these symbols as well. So naming conventions and a glossary of terms are part of a proper BPMN style guide in addition to modeling rules. There are some promising proposals, but their effectiveness needs to be established. So far, our evidence here is at Level III or below.

Actual Use of BPMN

Together with my colleague Jan Recker, I decided to pursue question three. We collected BPMN diagrams from consultants, students, and added diagrams that we found on the web, more than 120 diagrams in total. We kept track of which BPMN symbol was used in which diagram and plotted the resulting frequency diagram. The shape of this diagram was unmistakable: A long-tail distribution, with few symbols that were used very frequently (such as Activities and Normal Flow), and a number of symbols that we did not encounter at all (such as Compensation Transactions).^x Since we published our first results, a number of people have come forward and sent new diagrams to us, so this is an ongoing effort (and if you can share diagrams we much appreciate your participation).

We are interested in the practical use of BPMN for a number of reasons. Analyzing actual models tells us something about the typical level of complexity contained in a process model. This, in turn, could lead us to develop guidelines whether a model is significantly above or below the average level of complexity.

Looking at the symbols used in BPMN diagrams we can determine whether there is a common subset that modelers typically use. Our findings indicate that if there is such a subset it is very

small. The typical BPMN model we analyzed contained between 7 and 9 different BPMN constructs – That's not the overall number of graphical elements, but the number of *different* graphical elements – less than 20% of the overall BPMN vocabulary. We found anecdotal evidence that consultants use less constructs, while students use more – but our sample set was too small to determine whether this was statistically significant.

Our study has delivered a lot of insights that warrant further investigation, and the possible implications of the statistics have generated a lot of debate.^{xi} What we do know is that with a limited BPMN vocabulary it is possible to understand the majority of the BPMN diagrams created to date. This bodes well for communicating process models in organizations. We also know that there is no core subset of BPMN that the majority of modelers have agreed on. There were few similarities in the BPMN subsets used in different models – Most of them shared 3-4 constructs, but differed on the others. This indicates that modelers might choose a problem-specific subset, rather than work from a more general subset of the language. This may be a result of the different origins of our models. We are looking at models from a single organization to see if there is more consistency in the BPMN subset within one domain. Our evidence is currently between Level II-3 and II-2.

What we don't know

Our study tells us how people use BPMN, but it does not tell us why. For example, we saw between 7 and 9 different BPMN constructs in the average diagram. There are many possible explanations why a model may not contain more elements. It may be that the modeling purpose for each diagram did not require a more extensive use of BPMN. To establish this, we would have to correlate each diagram with its intended use (which we don't know). It may be that the modelers may not be aware of the meaning of the more technical symbols of BPMN. In order to find this out, we would have to administer a survey to the modelers on how much BPMN they actually know.

Our study is a snapshot of the average BPMN use to date, and our findings are limited to the symbols people actively use. There may be a significant difference between the BPMN that modelers understand (i.e., the symbols they can read) and the BPMN that they actually use (i.e., the symbols they apply). To assess the BPMN literacy we would have to administer a test, showing BPMN symbols to our audience and asking them to explain the meaning of these symbols.

It is possible that the tools used to create the diagrams did not support the entire BPMN vocabulary. To establish this, we would have to know the tool (and version) used to create each model – which, in many cases, we do, but in many others we do not. However, there is a relationship between the technology we use, and the things for which we use the technology. The capabilities of process modeling tools may shape the way people model processes, and if these capabilities change the way people model may change.

A proposition based on these observations would be that the BPMN vocabulary in active use will increase over time as the capabilities of the BPMN modeling tools improve. But this would also be the case if the BPMN literacy of the modelers improves. Or we might see the opposite effect: We found some evidence that novice modelers tend to use more language constructs than experienced modelers. Maybe this is a function of an effect such as “I spent so much time learning this language, now I want to show that I have mastered all of it.” It would be interesting to determine when modelers voluntarily restrict their use of a language and why.

It is notable that the most frequently used BPMN elements mirror those elements we find in other process modeling languages – most prominently, in flowcharts. This could be caused by modelers migrating from flowcharts to BPMN that essentially draws diagrams similar to their previous work using a new notation. Or it might indicate that the vocabulary of flowcharts is sufficient for many modeling purposes. However, BPMN provides a pathway to enhance these

diagrams with modeling constructs that support the implementation and automated execution of processes, while flowcharts do not. To what extent this capability will be used in practice, we do not know.

Our study tells us what the popular elements of BPMN are, but not whether the resulting models were good models. Work has been done on the automatic identification of errors in process models^{xii}, but these errors can be on the syntactic side (improper use of the language, i.e., modeling mistakes) and on the semantic side (improper representation of the process, i.e., wrong activity and decision labels). Syntactic errors can be detected with validation tools. Detecting semantic errors requires an understanding of the modeling domain; only case studies would help us in this regard.

So what's the Point?

Designers of modeling languages are engaged in anticipatory standard setting. They create a vocabulary and construction rules so users can employ the resulting language to create diagrams that the designers typically cannot foresee. This is the case with BPMN: It was designed as a universal process diagramming technique, allowing for process modeling both at a conceptual and at an implementation level. That's the intended use of the language.

Once a language specification has been publicly released, the user community takes over and the influence of the designers fades. The users may share the designers' intentions, but there is a chance that they may appropriate the language for a different purpose, mash it up with a different representation, or use it selectively without realizing its full potential. BPMN has a lot of potential. Understanding to what extent users tap into it helps designers assess the current understanding of their specification. It allows them to issue clarification and guidance to move users toward the intended use of the language.

Finally, trainers and educators need to align what their students know (current language use) with what they should know (recommended language use). The recommended use can be based on the intentions of the language designers, but the emergence of common design patterns, language dialects, and common areas of application should influence it. Gathering evidence about the actual use of BPMN allows us to determine both the starting point and direction of BPMN education efforts.

Collecting Evidence

Collecting evidence takes time. The most prominent example is the seminal work by Stephen R. Barley on the impact of new technologies on work.^{xiii} Barley started with two simple questions: What do technicians (as opposed to business people) do and what do they know? To study these questions, Barley observed – for an entire year – the interaction between radiologists and technicians in two Boston-area hospitals, both before and after the introduction of CT scanners. Both organizations were healthcare providers, introducing the same technology to perform the same function. Yet, in one hospital the technicians acquired skills from the radiologists to the extent that they could make routine decisions based on scan results, while in the other hospital the radiologists maintained a tight control over the CT scan procedure, marginalizing the role of the technicians. To validate and broaden his results, he performed a similar study again – ten years later.^{xiv} Scholars are still working through the implications of this landmark study.

For most businesses, 10 years are several lifetimes (and generations of executives). But in academia, there is time to work on those hard problems. Let's go and collect some evidence.

Michael zur Muehlen is Assistant Professor of Information Systems at Stevens Institute of Technology in Hoboken, NJ, where he directs the Center of Excellence in Business Process Innovation and coordinates the Graduate Program in Business Process Management and Service Innovation. Michael is a fellow of the Workflow Management Coalition and chairs the WfMC working group "Management and Audit." SAP Research, the US Army, the Australian Research

Council, and various companies have funded his research, which he has presented in more than 20 countries. He is the author of a book on workflow-based process controlling, numerous journal articles, book chapters, and conference and working papers on workflow and process management. Michael regularly advises corporations and BPM vendors on strategic and technology issues surrounding Business Process Management. He holds a PhD and an MS in Information Systems from the University of Muenster, Germany. He blogs at <http://www.bpm-research.com>

-
- i Special thanks to George Royce (University of Nebraska at Omaha) for providing this metaphor.
- ii The National Science Foundation has recently funded a multi-year project to investigate this area, led by Jeffrey V. Nickerson (Stevens Institute of Technology), James Corter and Barbara Tversky (Columbia University) (NSF-IIS 0725223). The results should be interesting.
- iii Knutilla, A., et. al., "Process Specification Language: Analysis of Existing Representations," NISTIR 6133, National Institute of Standards and Technology, Gaithersburg, MD (1998).
- iv German for: A song you heard in the morning that you cannot get out of your head. Don't ask why.
- v See for instance: White, S. (2004). Process Modeling Notations and Workflow Patterns. Workflow Handbook, 265-294.
Wohed, P., van, der Aalst, WMP, Dumas, M., ter Hofstede, AHM, & Russell, N. (2006). On the Suitability of BPMN for Business Process Modelling. Proceedings 4th International Conference on Business Process Management (BPM 2006), LNCS, Vienna, Austria.
- vi List, B. & Korherr, B. (2006). An evaluation of conceptual business process modelling languages. Proceedings of the 2006 ACM symposium on Applied computing, 1532-1539.
- vii See for instance: Rosemann, M., Recker, J., Indulska, M., & Green, P. (2006). A Study of the Evolution of the Representational Capabilities of Process Modeling Grammars. Advanced Information Systems Engineering-CAiSE, 4001, 447-461.
Wahl, T. & Sindre, G. An Analytical Evaluation of BPMN Using a Semiotic Quality Framework. *CAiSE'05 Workshops*, 1, 533-544.
- viii www.brsilver.com/wordpress/
- ix Recker, J., Rosemann, M., & von Uthmann, C. (2000). Guidelines of Business Process Modeling. vd Aalst, W., Desel, J., Oberweis, A.(eds.), Business Process Management: Models, Techniques, and Empirical Studies, Springer, Berlin et al, 30-49.
- x The complete results have been published here: zur Muehlen, M.; Recker, J.: How Much Language is Enough? Theoretical and Practical Use of the Business Process Modeling Notation. Proceedings of the 2008 Conference on Advanced Information Systems Engineering. Montpellier, France, June 16-18, 2008. Available at www.bpm-research.com/publications/papers.
- xi For a summary of the arguments (and counterarguments) see Sandy Kemsley's Blog: www.column2.com and Bruce Silver's blog: www.brsilver.com/wordpress/.
- xii Mendling, J., Moser, M., Neumann, G., Verbeek, H. M. W., van Dongen, BF, & van der Aalst, WMP. (2006). Faulty EPCs in the SAP Reference Model. Proceedings of BPM, 4102.
- xiii Barley, S. R. (1986) "Technology as an occasion for structuring: observations on CT scanners and the social order of radiology departments." *Administrative Science Quarterly*, 31:78 108.
- xiv Barley, S. R. (1996) "Technicians in the workplace: Ethnographic evidence for bringing work into organization studies." *Administrative Science Quarterly*, 41:404-441.