ON INNOVATION TIMING

Darrell Mann

Systematic Innovation Ltd, UK
E-mail: darrell.mann@systematic-innovation.com

Abstract

This paper considers the complex issue of innovation timing. Building on an extensive programme of research and application evidence, the paper outlines the key factors that affect and ultimately determine *when* innovations are or are not likely to succeed. A prototype timing algorithm is presented. The algorithm distinguishes three main innovation timing answers; 'now', 'never' and 'at some point in the future'.

1. Introduction

The technology discontinuous-evolution trend tools of classical TRIZ are known to offer unrivalled means by which future generations of products and process might be predicted. They are, however, relatively weak when it comes to answering questions concerning *when* a particular innovative step is or is not likely to succeed in the market. This issue of innovation timing is vitally important to any organisation. It is at the same time an extremely complex question to try and answer due to the myriad different factors that can influence the final success-or-fail outcome.

The aim of this paper is to present a high level perspective of the factors that will influence innovation timing, and to then attempt to distil from them some of the issues that are more significant than others. As such, the paper intends to build on the initial discussion on innovation timing found in the Hands-On Systematic Innovation book (Mann, 2002).

The paper emerges following a prolonged programme of research and live project work with clients. This work has had two primary aims. First, an uncovering of the underlying innovation dynamics governing which timing factors are more important than others and in what situations. Second – and perhaps more importantly – then, has been the testing and verification of findings on real life innovation timing situations. The focus of this paper is on the first of those two parts of the innovation timing story.

A significant portion of the paper is devoted to detailing what have been determined to be the most important innovation timing factors. The factors included in this analysis are:

1) emerging conflicts between linear market trend directions
2) business model viability – where the essential elements required for any innovation to succeed are outlined, and their significance to timing issues discussed
3) whether the innovation is incrementally-enhancing, discontinuously-enhancing, or discontinuously-inferior (i.e. 'disruptive')

4) whether or not there are 'hidden' failures in an existing market
5) whether there are symbio-genesis opportunities that can be created and exploited
6) how dynamic the market sector has been in the recent past

The paper describes tools and strategies that have been developed and can be used to derive the essential innovation timing knowledge for each factor. In so doing, the paper demonstrates that no single factor alone can offer a sufficiently robust innovation timing model. A final section of the paper seeks to demonstrate that certain of the factors can, however, be combined into a 'timing algorithm' that, although inevitably at an early stage of its evolution, can still be practically deployed to give meaningful insights into a wide range of specific innovation situations.

The discussion begins with an examination of the last of the above described innovation timing factors – how dynamic the market sector has been in the past – since this is closest conceptually to 'traditional' means of market forecasting.

2. Market Dynamics

Given the importance of innovation timing, it is not surprising that there are many tools and techniques designed to help organisations to predict the future. Without exception, they have all ultimately failed to offer any kind of meaningful insight into the dynamics of the evolution processes. One of the main reasons why such tools and techniques fail so badly is that they are essentially built on an assumption that the world operates in a linearly predictable manner. The simple truth, however, is that the world does not work in this way. In line with the now clichéd expression involving the butterfly flapping its wings in Brazil creating a tornado half a world away, apparently minor shifts in one thing can lead to some highly non-linear outcomes somewhere else. Paradoxically, almost none of the myriad economists, mathematicians, psychologists and anthropologists working in the field appear to actually believe the linear assumption. Unfortunately they continue to try and build ever more sophisticated mathematical models to try and compensate for the failings in their existing models.

Looking back at the evolution of the trends part of TRIZ, we may see an implicit assumption that even though it may not be possible to construct a mathematical model of the non-linear dynamics of the world, it is highly possible to identify what the non-linearities *look like*. Crudely speaking then, the TRIZ trends may be seen as maps of discontinuous evolution; each jump along one of the trends in essence representing the shift of a system from one way of doing things to another. A shift, in other words, from one s-curve to another s-curve.

The models traditionally used by forecasters are built on historical data. Moreover, it is usually historical data which is made to fit to some mathematical algorithm. Once modelled, the basic idea is then that we extrapolate from the known history into the future. The inevitable non-linearities soon make such extrapolations meaningless. The Evolution Potential concept described in (Mann, 2002) and elsewhere can likewise be used to build a historical picture of an industry. Beyond that connection to historical data, however, it is profoundly different. It is different because it is fundamentally about mapping non-linearities. Each and every jump along one of the trend spokes featured in the radar plots, in other words, represents a non-linear jump that has taken place in that particular system.

What this means is that if we plot the Evolution Potential of a given system (either technical or business) at one point in time, and then again at another time, it is possible to determine the *discontinuity rate* of that system. Figure 1, for example, illustrates a hypothetical example showing how the plots 'flower' over time, as progressively more and more of the available potential in a given system is exploited.
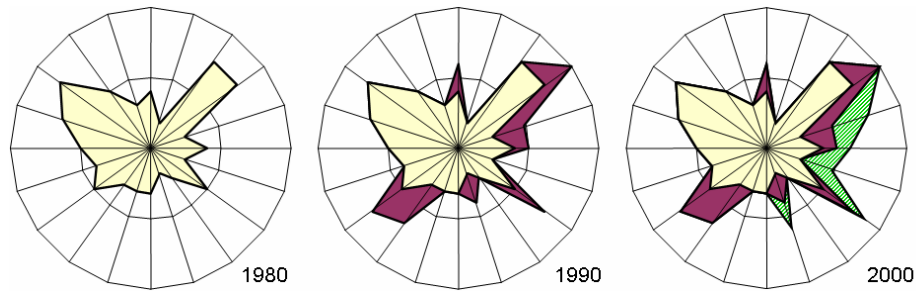
Figure 1: Evolution Potential Plot 'Flowering' As A Function Of Time Gives 'Discontinuity Rate'

In the experience of the author in drawing several tens of thousands of such radar plots, they are probably as good a future-prediction tool as is possible based on historical data. By way of example, (Mann, 2004a) and (Mann, 2005) describe a pair of related industry evolution analyses – one on fan technology, and the other on compressors – in which it is possible to see the discontinuity rate (i.e. rate at which the plot 'flowers', usually measured in number-of-discontinuous-trend-jumps-per-unit-time) remaining almost constant throughout the evolutionary history of those systems. This 'discontinuity rate' concept will be the subject of a second article to follow this one in the next issue of TRIZ Journal (Mann, 2006).

Just because discontinuity rate has been constant in the past, of course, does not mean that it will necessarily remain so in the future. In our research, we believe that it is 'likely' that it will, but on the other hand 'likely' and 'certain' are two very different matters.

Overall conclusion: forecasting based on discontinuity rate can be useful; the data is a useful piece in the overall timing equation, but don't rely on it exclusively.

So much for innovation timing analysis based on historical data. We will now switch our attention to factors based on forward rather than backwards looking:

3. Business Model Viability

We start on these other factors by looking at another part of the classical TRIZ toolkit; this time the Law of System Completeness. Whilst usually applied to technical systems, the 'Law' may be seen to be equally applicable to business systems in general and any system we may care to attach an 'innovation' label to in particular. What the Law of System Completeness tells users is that there are a certain number of elements that will need to be present in a system in order for it to be 'viable' (Mann, 2004b). In technical terms, these elements comprise Engine, Transmission, Tool, Interface and Control. In the terms of an 'innovation system' we might interpret these five elements as shown in Figure 2.

If any one of these elements is missing, the Law tells us, then the issue of innovation timing becomes somewhat moot. The issue in such situations is about establishing the missing elements in order to achieve the minimum viability requirements and not about timing. One of the elements, 'route to market' (usually interpreted as 'transmission' in the classical model of the Law) is frequently the one that is missing from the system. This is particularly so for disruptive innovations and for innovations derived by small, independent new players. This is so because it is usually the incumbent – who has already established a route to market – is likely to prevent a threatening newcomer from exploiting such routes.
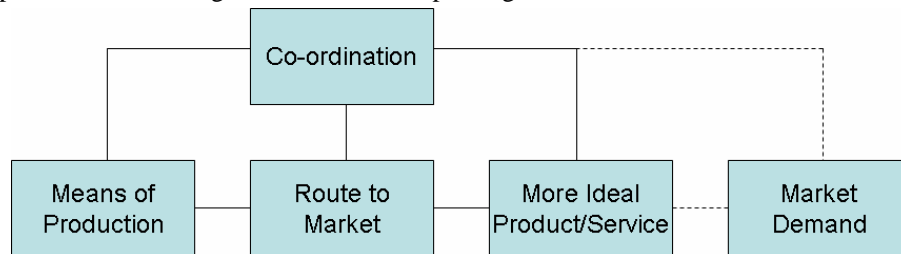
Figure 2: Essential Elements Of 'Innovation System'

With respect to the issue of the often difficult market demand, it is also worth noting a need to distinguish between technology and market. We may do what is necessary quite simply by constructing a 2x2 matrix describing new and existing technologies and markets as shown in Figure 3.
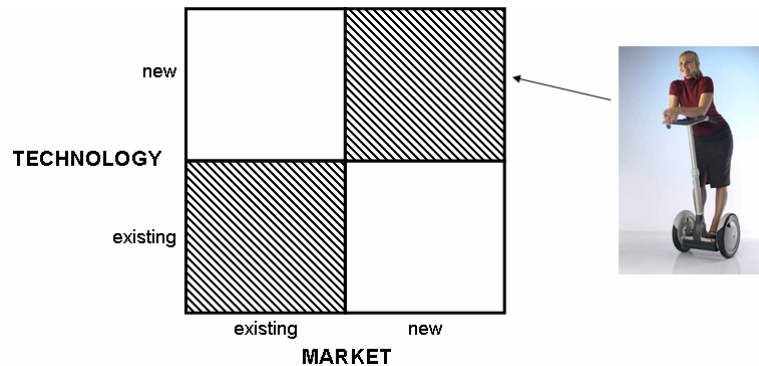
**Figure 3: Innovation Arena – Technology And Market**

The Figure 3 matrix presents another yes/no gate in the innovation timing question. A somewhat different perspective on this figure may be found in (Monnier, 2004). In Monnier's article, the thinking behind whether to play in the new-technology-and-new-market box is inconsistent with our findings, and so may necessitate some research to understand the differences. The question accompanying our proposed yes/no decision is whether the innovation in question may be seen to reside in either of the shaded boxes. If it does, then the likelihood of the innovation being successfully deployed is negligible. In the existing- technology/existing-market box, the innovation is likely to fail either because customers are likely to react with a 'so-what?' attitude or because the provider is forced down a 'low-cost' route which

inherently puts pressure on profit margins. Perhaps less obviously, innovations sitting in the new-technology/new-market box are also unlikely to succeed. The Segway, featured in the Figure is a good example of this kind of innovation. Until such times as consumers can connect one of either the technology or application to something they are already familiar with they will not accept the innovation. Within this constraint, if history repeats, the Segway remains destined to reside in the category marked 'interesting, but no thanks'.

So, taken all together, the issue of business model viability is best interpreted as a series of discrete yes/no gates. Failure to pass successfully through all of the gates means that the answer to the innovation timing question is 'never'.

4. Innovation Type

Another discrete yes/no type factor in the innovation timing process concerns the issue of the type of innovation. Our thinking in this area is primarily based on the work of Clayton Christensen (Christensen, 1997), and in particular his ideas on incremental, breakthrough and disruptive innovation. 'Disruptive' in Christensen's definition involves an innovation which is seen by the present incumbents and customers as inferior. Disruptive innovation is important to the understanding of innovation timing issues thanks to Christensen's identification of the so-called 'innovator's dilemma'. As illustrated in Figure 5, what Christensen identified is the mis-match between advances in technology and customer expectations. The innovator's dilemma exists because the slope on the technology line is steeper than that of the customer expectation. Hence a disruptive innovation, although it may be 'inferior' in its initial form, will eventually evolve to a point where it catches up with the customer expectations of the incumbent, eventually then taking over the market from the initial incumbent.
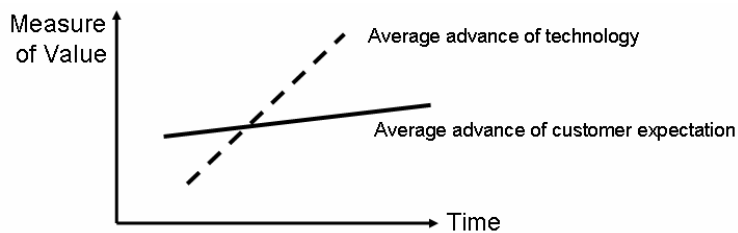

Figure 5: Classical Christensen Innovator's Dilemma

The first yes/no gate relating to Christensen's model relates to the breakthrough innovation category whereby, if the breakthrough innovation is derived by a small, independent new player, it is very *un*likely to succeed. In some ways this yes/no rule is closely linked to the earlier discussion on business model viability. On the other hand, it is considered different enough to merit separate attention here.

Taking this idea of 'breakthrough' a step further, it is worth noting this author's revised definition of the term (Systematic Innovation, 2005c). In this model we use the Evolution Potential concept to compare an innovative solution against the incumbent. A genuine breakthrough is then observed if, when these two plots are overlaid on top of each other, the evolutionary state of the innovation is the same or more advanced along *all* of the spokes on the plot. In situations where an innovation

is more advanced along some lines, and has passed backwards along other trends – as in the field-based keyboard innovation illustrated in Figure 6 – then it does not count as a genuine breakthrough.

The market failure of the light-keyboard provides a clear illustration of another yes/no in the innovation timing story: namely if the innovation is not a genuine breakthrough (as defined above) then customers are being presented with one or more trade-off decisions, any one of which may prevent the possibility that the innovation will succeed. A genuine breakthrough that also satisfies the business model viability criteria, on the other hand, is highly likely to give an innovation timing answer of 'now'. In any other scenario, the timing issue requires a deal more analysis.
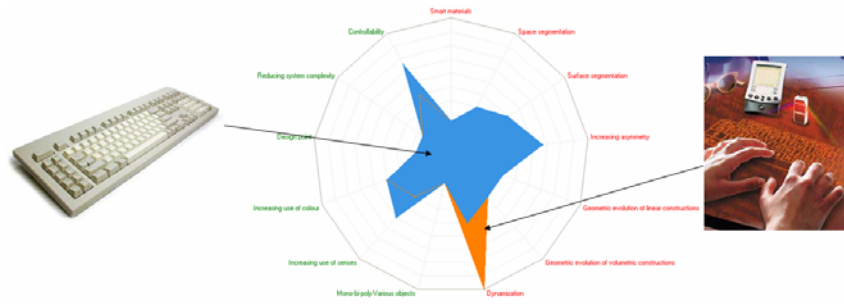
Figure 6: Use Of Evolution Potential Plots To Determine 'Breakthrough' Status

Christensen's model offers additional insights in situations where 'more analysis' is required, albeit in a modified form. The Christensen model shown in Figure 5 assumes that the technology advance and customer expectation advances are linear. The first correction that needs to be applied, then, is to recognise that the technology advance characteristic is actually an s-curve (Mann, 2002) rather than a straight line. The second correction involves the customer expectation evolution. Here we need to get into a whole new level of complexity if we are to make any sense of the innovation timing question:
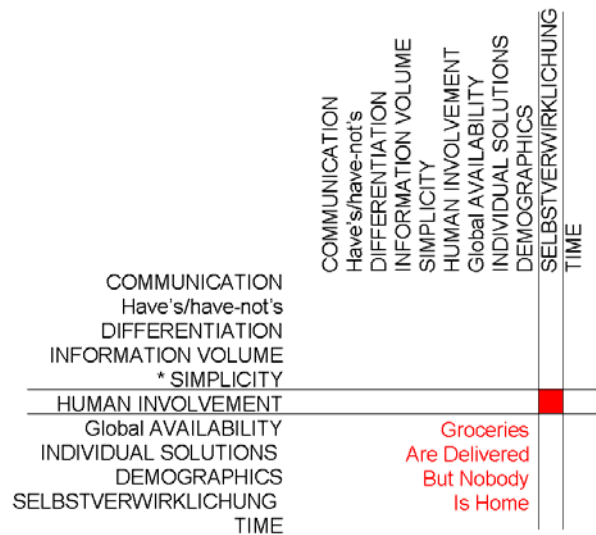
5. Trend Conflicts

Here is where the innovation timing question reaches the peak of its complexity. And also the least mature part of the timing methodology. The focus here shifts to continuous (as opposed to the TRIZ 'discontinuous') consumer and business trends. By 'continuous' here, we mean things like demographic trends, and the sorts of things to be found in publications like Megatrends (Naisbitt, 2000) or 'Dictionary of the Future' (Popcorn, 2002). The key difference, however, between the traditional use of such trends and what is necessary in order to understand innovation timing issues is that the trends need to be considered in combination rather than – as is traditional - individually. Consider any one of these trends individually and it is usually possible to see the logic and validity of that trend. Things start to go wrong – i.e. the trends don't work – because sooner or later the advance along one trend direction conflicts with an advance along another trend.

What is required, then, in order to make sense of these trends and the ways in which they interact is to construct some kind of conflict analysis tool. For this author,

this means constructing a trend conflict matrix of the sort illustrated in Figure 6. The idea here is relatively simple, and usually involves the following steps:

1) identify the trend directions that are relevant to the innovation under investigation
2) use the Perception Mapping tool (Mann, 2004b) in order to map the relations between these trends and from there determine which trends are more important than others
3) for the most important trends, construct a trend matrix of the sort illustrated in Figure 7, identifying for each box any conflicts that may emerge between the pair of trends that define that box

The job of constructing and populating such a matrix is beyond the scope of this paper. The level of effort required to create and populate the matrix is highly dependent on the number of trends that are to be included in the analysis. Typically, meaningful answers can be derived by considering between 15 and 20 trends, but it is too early in the evolution of the concept to be able to deliver any concrete guidelines.

Figure 7: Consumer Trend Conflict Matrix Structure

The best outcome that may emerge from this trend conflict analysis is that there are no significant conflicts present during the likely timescales of the innovation under investigation. Should this be the case, then the calculation of innovation timing becomes a relatively simple one. Should there be conflicts, however, then the innovation timing question continues to be a highly complex one. Before we explore exactly how complex, it is necessary to examine two final factors that can have a potentially significant impact on the innovation timing question:

6. 'Hidden Failures'

A 'hidden failure' is a situation in which an incumbent product or service fails to meet one or more customer needs, but in such a way that the customer is unable to either see or describe that failure. The usual examples of hidden failures are the pizza-

box that opens to reveal a cold, soggy pizza and the digital camera. In the first case, customers experiencing their first cold, soggy home delivery pizza are conditioned to expect that all home-delivery pizzas are the same. In the case of the camera, the emergence of the digital variant exposed the hidden failure of film cameras. The hidden failure in the film camera is the need to process pictures off-line. It is not until someone emerges with an innovation that eliminates the hidden failure that customers become aware that it was present.

Being the first to identify and resolve a hidden failure is an excellent means of creating innovative solutions where the answer to the timing question is 'now': Find a hidden failure in an incumbent product or service and your innovation will find an audience that is immediately receptive.

That is the good news part of the hidden failures story. The bad news is that finding this type of failure is not easy – surprise, surprise; if they weren't hidden, everyone would be able to travel this route to innovation success. Several of the TRIZ trends can help to reveal hidden failures if used in a way we describe as the 'voice of the product'. What this means is that the trends of discontinuous technology evolution are designed to act as signposts to good solutions. If we can then identify why a particular trend direction is useful in our specific instance, then we are well on the way to locating hidden failures. More recently (Systematic Innovation, 2005a), we have described a more sophisticated means of identifying and capturing the hidden failures.

7. Symbio-Genesis

Less reliable, but nevertheless equally useful in helping to unravel the innovation timing question is the idea of symbiogenesis (Margulis, 2003). In many ways similar to the TRIZ Mono-Bi-Poly trend, what Margulis reveals in her work is that one of the most frequent and most reliable form of speciation (the natural world equivalent of innovation) occurs when existing organisms decide to merge in order to achieve a symbiotic benefit. The classic win-win innovation. We see a similar record of innovation success in the unrelated work of Metcalf (2002), on the emergence of new words into the dictionary. (Systematic Innovation, 2005b) further analyses Metcalf's work to reveal that close to two-thirds of all new words entering the language emerge from symbiogenetic roots.

We may observe many examples of symbiogenesis (or advances along the Mono-Bi-Poly trend) in the commercial innovation environment too. In fact, as in nature and etymology, the symbiotic merger of things that have been successful individually is an extremely reliable indicator of innovation success. Typical examples include recent mergers of things like toothpastes into chewing gum, Virgin Brides (combining all of the individual wedding activities into a one-stop shop), and financial institutions combining with house-builders to offer easy home purchase schemes. Providing the merger is a truly symbiotic one (i.e. no negative jumps along any of the TRIZ trends), then from an innovation timing perspective, the answer is very highly likely to be 'now'.

8. An Innovation Timing 'Algorithm'

The preceding discussions reveal a not-unexpected high level of complexity in relation to the innovation timing question. A helicopter-level review of the different factors considered reveals three basic possible timing outcomes:

      1) the innovation will never succeed and so timing is irrelevant

2) the right time to introduce the innovation is 'now', and
3) the innovation is likely to be successful if launched at some point in the future.

Many have attempted the feat of predicting innovation timing, and none have got it right. The likelihood, therefore, of overturning that poor historical record in one short paper is pretty slim. What we can hope to achieve, however, is some level of prediction reliability in a variety of different circumstances. The aim of this section is to distil from the preceding discussions, rules and guidelines that may be applied to the innovation timing question. The easiest way to accomplish this is by examining each of the three timing outcomes separately:

8.1 Never
Innovations will not succeed (and hence the timing issue is irrelevant) if:
* any one of the conditions of business model viability are not satisfied
* a genuine breakthrough innovation is derived by a small, independent organisation, and that organisation seeks to compete with an established player directly

8.2 Now
Innovations will succeed immediately if all of the conditions of business model viability are satisfied *and*:
* a 'hidden failure' is identified and exploited, or
* a genuine breakthrough innovation is created or acquired by a large player, or
* a symbiogenesis opportunity is identified and exploited

8.3 At some point in the future
This is clearly the most difficult of the three possible timing scenarios as it requires some kind of calculation to be performed. Firstly, of course, the tests of business model viability must have been passed. Beyond that, there are a number of alternative scenarios that may arise which will influence how to calculate innovation timing. These are:
* if the innovation is not a genuine breakthrough (i.e. there are technical trade -offs relative to an incumbent solution) the timing will be correlated to the level and type of trade-offs present
* if there are conflicts emerging from different customer trends that already exist or can be expected to emerge in the desired lifetime of the innovation, the timing will be correlated to these trend conflicts (and to the above technical trade-offs)
* if the innovation is incremental in nature then timing will be correlated to Position on s-curve and prevailing customer expectations

Each of these 'if' statements will necessitate an innovation timing 'calculation' that will, for example, require the construction of a trend conflict matrix. As may be expected, such activities are both time consuming and highly specific to each individual case. It is not possible for us to enter a meaningful dialogue on such calculations in this paper. What we can say, however, is that the way in which we utilise the information from these 'if' calculations is as modifiers to the default timing calculation process.

We can begin to draw together a calculation process if we define a number of correction factors. One correction factor is needed for each of the three 'ifs': F_b

representing a calculation modifier that can be calculated for situations where there is no genuine breakthrough; F_t representing a similar modifier for instances where there are trend conflicts, and F_s representing a modifier for the third 'if' case where there is an incremental innovation and we need to account for position on s-curve,.

Before incorporating these F factors, if none of the 'ifs' is present, then the innovation timing calculation is made solely as a function of the discontinuity rate calculated from the evolution potential analysis described earlier in section 2, 'Market Dynamics'. In this situation:

$$\text{Innovation Timing} = f\{\text{discontinuity rate}\}$$

Presence of any of the 'if' elements then modifies this calculation as follows:

$$\text{Innovation Timing} = f\{\text{discontinuity rate}, F_b, F_t, F_s\}$$

The whole innovation timing calculation process can then be summarised as shown in Figure 8:
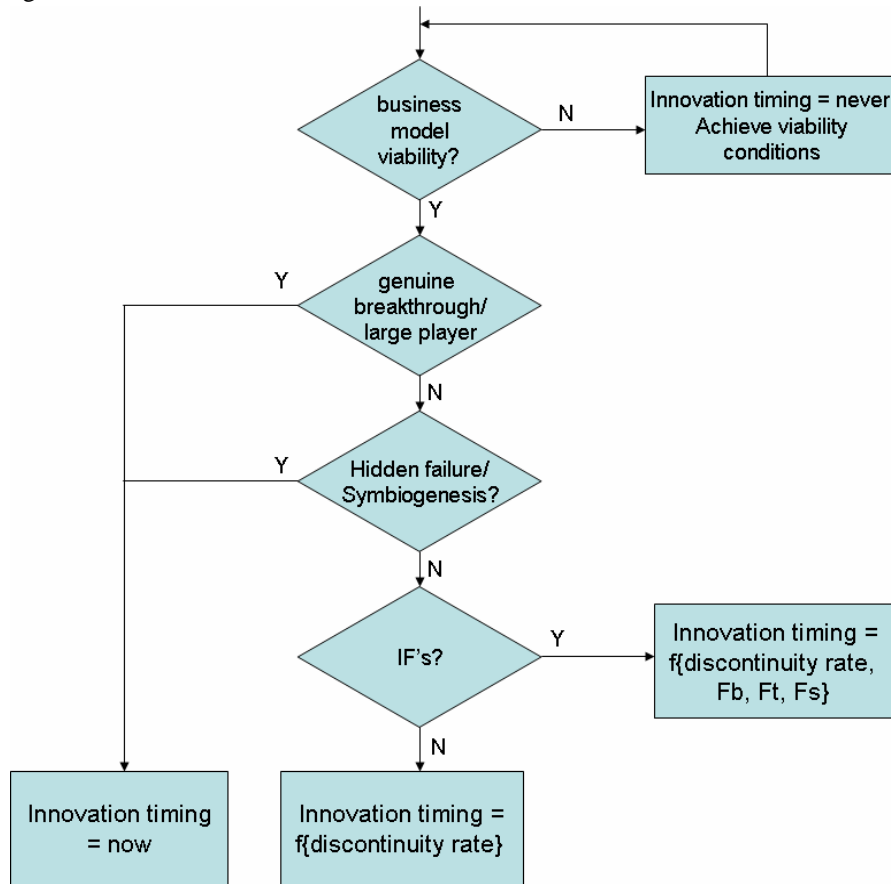
Figure 8: Overall Innovation Timing Calculation Process

In essence, then, what we see here is a series of discrete yes/no factors that pointing to a now or never timing decision, and the use of 'discontinuity rate' (calculated from historical data) to establish the in-between timing possibilities. This discontinuity rate calculation is in turn modified according to the existence or otherwise of a number of other discrete innovation timing factors.

9. Final Thoughts And Next Steps

TRIZ is an extremely potent tool in terms of defining the 'what's' of innovation. The discussion here has centred around the issue of 'when', and we have hopefully revealed not only the enormous complexity of the problem, but, more usefully, a few repeatable rules and calculation structures that can be used to answer the timing issue. There is still much more data to be collected and analysed before we can be confident that the true underlying dynamics of innovation timing have been captured and encapsulated. Clearly this is work that will continue in our research activities into the innovation process. We also, however, clearly understand the issue of 'it's the whole thing, stupid' in any of the work we do. Thus, in addition to what and when, we also need to create a clearer picture of who, where and why. The task of understanding *when* to launch an innovation onto the market is inevitably coupled (we believe) to the other questions. That coupling means that what might be the right timing decision for one geographical region or customer niche may be entirely wrong for another sector. Unfortunately, it is our belief that we cannot get to the 'whole story' without first gaining a better understanding of the individual questions. We believe that 'when' is the most difficult of the questions, and that as we start to incorporate findings on who, where and why, the timing factors discussed here will continue to be relevant. We may find other factors, and we may find that the relative importance of each needs to be shifted, but we also believe that the basic innovation timing framework discussed in this paper will continue to be applicable. The trick – if there can be such a thing – during the research we believe is to simultaneously keep an eye on both the micro-detail and the bigger picture. To be both zoomed-in and zoomed-out at the same time.

The aim of this paper has been to capture the methodological essence of the innovation timing issue. It is our intention to publish papers providing case-study examples showing how the theory translates to practical reality.

References

Christensen, C., (1997), 'The Innovator's Dilemma', Harvard Business School Press.
Mann, D.L., (2002), 'Hands-On Systematic Innovation', CREAX Press.
Mann, D.L., (2004a), 'Fan Technology: Evolution Potential And Evolutionary Limits', paper presented at IMechE International Conference On Fans, London, 9-10 November, 2004.
Mann, D.L., (2004b), 'Hands-On Systematic Innovation For Business And Management', IFR Press, 2004.
Mann, D.L., (2005), 'Compressor System Technology: Evolution Potential And Evolutionary Limits', paper presented at IMechE International Conference On Compressors And Their Systems, London, 4-7 September, 2005.
Margulis, L., Sagan, D., (2003), 'Acquiring Genomes: A Theory Of The Origin Of Species', Basic Books.
Metcalf, A., (2002), 'Predicting New Words', Houghton Mifflin Company, New York.
Monnier, B., (2004), "Application of the TRIZ Method to Business Management Activities", paper presented at ETRIA TRIZ Future 2004 conference, Florence, Italy, November 2004.

Naisbitt, J., Aburdene, P., (1991), 'Megatrends 2000', Pan Books, London.

Popcorn, F., Hanft, A., (2002), 'Dictionary of the Future: The Words, Terms and Trends That Define the Way We'll Live, Work and Talk', Hyperion.

Systematic Innovation e-zine, (2005a), 'Finding The Missing Functions', Issue 34, January 2005.

Systematic Innovation e-zine, (2005b), 'Best of The Month – Predicting New Words', Issue 41, August 2005.

Systematic Innovation e-zine, (2005c), 'Defining Breakthrough Solutions', Issue 41, August 2005.