



An Introduction to

# INFORMATION NORMALIZATION THEORY

Humour's role in the war against infected memes

ALASTAIR CLARKE

Copyright Dot Net Press 2010

All rights of the author reserved

Reproduction of this publication is strictly prohibited. However, educational institutions may print or circulate the material herein to their students provided it remains in its current form and undergoes no alteration, and no charge is made for its distribution or reproduction. If such educational circulation occurs we request only a courtesy confirmation of numbers to [info@alastairclarke.net](mailto:info@alastairclarke.net). This article may not be reproduced in any other publication or similar venture, including the web, without prior permission from the publisher.

[www.alastairclarke.net](http://www.alastairclarke.net)

# AN INTRODUCTION TO INFORMATION NORMALIZATION THEORY

One particular aspect of human life sets our species apart from all others. It is not the way in which we function as a society, care for the weak, wage war on each other or farm the land. Neither does it equate to our education of the young, the desire to explore and conquer, or our perennial resourcefulness through good times as well as bad. Yet it is in every one of these, and much of our remaining behaviour besides.

All species need to know how to behave. Every individual organism, regardless of how basic its physiology or intelligence, requires instructions for living. For nonhuman species, the majority or even all of the data providing these essential facts is inherited genetically. Photosynthesizing phytoplankton thrive where the light is strong because their systems are programmed to do so; many species of dormice conserve energy by hibernating for the same reason, and kittens play instinctively because it helps them to develop their predatory skills. These behavioural traits occur as if spontaneously, and each generation of the species engages in similar activities to the one before. During its lifetime, a conscious organism may employ its memory to help adjust its behaviour according to experience, but can only do so within certain parameters imposed by the underlying genetic code. With minor exceptions, alterations in this inherited information are painstakingly slow and limited in scope, and little happens that is not prescribed by its detail.

Humans, on the other hand, are different. The species still retains its genetically encoded instructions that direct and inform innumerable activities such as sleeping and waking or breathing and eating, but there is also a significant, second source of guidance available to the individual. From birth to death we receive more of our behavioural instruction via cultural means than any other species, and it is this ability to process and

transmit a constantly changing stream of data, this informational flux in which we live, that sets us aside from all other animals. In addition to tried and trusted genes being bequeathed from one generation to the next as partially shuffled hand-me-downs, *memes* can be spread swiftly from any person to any other in either direction, whether from old to young, young to old, or peer to peer, without ever being absorbed into the biochemical make-up of the species. A large proportion of our behaviour in the modern world, from how to put on clothes in the morning to the operation of a supercomputer at work, is learnt via this form of instruction. We receive information from another human being who has often also learnt it from elsewhere, we absorb it and process it and, if circumstances turn out appropriately, we later refer to it to guide our behaviour. And as we receive one bit, we just as frequently transmit another, which we too have either learnt from someone else or originated ourselves. For the modern human, culture is a fast-paced trading system with a continual round of exchange, and information is its currency.

While there are significant upsides to revising one's behaviour according to rapidly updatable programmes that can be selected or deleted at will, this ubiquity of data in human life comes with its pitfalls. The individual is primed to take instructions on board, genetically disposed to learn from experience and to track down and process new memes wherever they can be found. Failure to do so could mean missing out on cultural developments and falling behind society's expectations, and this, in turn, could lead to serious consequences for the individual's chances of selection as a mate or a peer. The sticking point is that learning is only as good as the information that is learnt, and this willingness to absorb data renders the individual vulnerable to that which is misleading, just as opening the gates to allow in the welcome leaves them open to undesirables too. In the world of cultural information exchange, these unappealing units are often innocently and inadvertently engendered by the occurrence of error, but a system of this kind also remains dangerously open to abuse by competitors. Human beings are not only the most learning of species, but the most deceiving as well.

Just as my erroneous directions may send you over the edge of a cliff, so may I intentionally mislead you towards the same messy end. Compounding this, since I have used the information exchange to instruct you rather than taking you there myself, I can be elsewhere as you plummet to your death, bewailing your sad demise to your family and friends, and painting in sombre colours the many ways I tried to save you as I worm my way into

the affections of your mate. In a competitive environment the employment of untruths is a significant weapon to acquire for your arsenal, and those who are more convincing in their dissemination of misinformation will reap the greatest benefits of this strategy. Your only defence is to manage somehow not to be taken in by my deception, enabling you to remain safe from the memetic infection that my misdirection would otherwise have caused. Unfortunately, there is often a crucial problem with such attempts at immunity. In much misinformation things are not quite as clear-cut as we might wish, and elements or factors of the misleading data or the manner in which it is presented may appear to be sound or convincing, making its rejection less straightforward. Here the individual is vulnerable to the *attractive error* with which they are confronted, and action must be taken to ensure the infected data is not adopted as a basis for future behaviour.

How, then, do we resolve the problem of wanting to absorb as much information as possible but needing to protect ourselves from that which might infect our minds and damage our chances of survival? According to information normalization theory the answer lies in the evolutionary adaptation known as humour. If this is so, the faculty performs a function that is much removed from the popular perception of the fun it provides. While its defence system against corrupt or malicious memes generates a great deal of pleasure, this amusement exists solely to motivate us to exercise a degree of circumspection towards all information. This is not to say we should underestimate the significance of these rewards, however, since the endorphin-fuelled highs of the humorous response are among the most frequent and intense of any enjoyed by human beings, which tells us a great deal about the importance of the underlying behaviour they reward. Woody Allen was not far off the mark when he commented:

“Sex is the most fun you can have without laughing.”

(From *Annie Hall*, among others)

But perhaps the stock witticism that *laughter is the most fun you can have with your clothes on* is even closer to the truth. For the continuity of the species the occurrence of reproduction is essential, which explains why sex is so enjoyable, yet here we have a pleasure motivating some other form of behaviour that is sufficiently significant to be considered in the same breath. Admittedly, most people would not pretend that laughter is more enjoyable than an orgasm, yet it occurs with a higher frequency and constitutes the most intense reward associated with cerebral activity.

This assessment of evolutionary value according to the intensity of the pleasure the organism receives tells us that the job that humour has done, the defence it has provided in the war against infected memes, has clearly proven mightily important to the human race.

So how does laughter manage to protect us from unreliable information? The mechanism by which the brain calculates the intensity of amusement the individual experiences harbours the answer, but to examine this central engine of humour we need first to define more closely the nature of the problem the faculty has evolved to address. As we know, unconvincing misinformation is not a threat to anyone, so humour does not need to compensate for it. Neither are we concerned about convincing data that qualifies as sound because this can only be good for the meme pool. This leaves us with two further combinations of these elements: unconvincing data that is correct, and convincing misinformation. The former may lead to a missed opportunity, and could therefore be considered wasteful, but in the majority of situations, barring useful warnings that the individual fails to heed, little harm will arise from its existence. The latter combination of convincing misinformation, however, could prove seriously dangerous to any who are not on their guard, and will rarely herald beneficial results unless the individual turns out to have been incorrect in their assessment of what constitutes erroneous data to begin with, or an accident of causality allows a chance advantage to develop from an unpromising foundation. According to information normalization theory, defence against such threatening data is achieved by offering rewards for the identification of its two distinctive but problematic features in one location: the presence of misinformation in its various forms and a subjective vulnerability to taking it seriously. When these factors coincide, the individual is in danger unless the misleading data can be identified as such and dismissed accordingly, and the more convincing we find it, the greater the threat. It is here that humour steps in to compensate, encouraging the individual to see through the veils of misinformation, to pick them apart and identify the memetic flaws within. If we successfully recognize an instance of misinformation for what it is, we have found the hidden pitfall beneath a beguiling surface, and receive a reward for our efforts. This vital balance of factors is described by the equation below, which summarizes the central activity of the system:

$$h = m \times s$$

I am willing to accept that this might seem a little drab for those not accustomed to dealing with humour in such terms, and a world apart from the warmth and enjoyment the faculty so often provides. We should not be surprised by this. The study of humour is sometimes one of contrary impulses, since the source material is capable of proving highly distracting from the seriousness of the matter in hand. However, the mechanism of humour is no less a product of biochemical processes than the mechanisms of digestion or respiration, or indeed of any source of physical pleasure such as orgasm, and deserves to be treated on an equal footing. Regardless of this defence of analyzing the system in this manner, I am not suggesting that the individual engages in an algebraic calculation before they laugh, only that this schematic interpretation represents the forces that pull the brain's responses to data in different directions.

Nevertheless, we need more information to make sense of this arithmetical interpretation of the most pleasurable of cognitive faculties. The result  $h$  is the humorous value we attach to the bit of data in question, also equating to the initial extent of the reward to be released to the individual and, all other factors being equal, the intensity of amusement that is experienced. I say 'initial' because secondary influences may increase or decrease this total elsewhere in the system, but we can afford to overlook these embellishments of the faculty for the purposes of this introduction. As a value, however,  $h$  exists not only in terms of the amused person's pleasure, but also as a summary of the risk that has been averted by the activity of the faculty. Thanks to the individual being encouraged to look more closely through the misinformation with which they were confronted, this particular pitfall will not now present the problems it might otherwise have caused. The two values of pleasure and averted risk are, of course, equivalent, since the greater the disinfection of hazardous material that has occurred, the greater the amusement with which the individual should be rewarded.

The remaining two variables in the equation constitute the major components of humour's engine. If both of these elements present positive values, so too will  $h$ , and we have a humorous response. The first,  $m$ , reflects the extent to which the information has been distorted or skewed from its default position. *How much misinformation is there?* is the question the brain is asking at this point. Clearly this is not an objective assessment that can tell us whether the incoming data is factually correct or not, but a subjective analysis by the individual depending on their knowledge and perception. What one person recognizes as misinformation may be seen as

entirely accurate by the next, and these differences in opinion will directly and profoundly affect the result produced here, and therefore whether the faculty is drawn into a response or not. You may consider my choice of clothing hilarious, for example, despite my interpretation of it as the height of sartorial expression.

Next, the second variable,  $s$ , represents the individual's susceptibility to taking seriously the misleading data that appeared as  $m$ . An alternative way of expressing this vulnerability is the *memetic value* the individual awards the perceived event, which is again the result of a purely subjective assessment.  $S$  is therefore the product of all the factors that influence the individual's decision about whether to be convinced by the information with which they are presented. *What associated and contextual metadata accompanies this misinformation that makes it appear worth looking at?* is the question the brain is asking at this point. For example, you may tell me you love my new shirt even though I know you do not. While your words constitute incorrect information to be processed and discarded, your tone of voice is convincing, thereby providing good elements in among the bad and increasing your comment's memetic value. Your lie constitutes the misinformation I have identified, but your voice represents a contextual factor that influences upwards my assessment of its value as a piece of information. The more convincing you make the delivery, the more likely I am to take you seriously, which increases my susceptibility to your lie which, in turn, gains a greater chance of survival as a successful meme in the cultural exchange. There are various kinds of misinformation and many factors that can affect this determination of memetic value, but the details are less important than the central principle that if you perceive something to be wrong then it *is* wrong as far as the faculty of humour is concerned, and if you find elements of the stimulus attractive or convincing you will inevitably find yourself susceptible to its misdirection to some extent.<sup>1</sup>

Another important point to note regarding misinformation is that it does not only reside in knowledge. Because such a large proportion of our instructional inheritance occurs culturally, the wide-eyed human is continually absorbing and processing the activity of others in case their behaviour constitutes a valuable lesson to be learnt. Purely because information is not passed on with the intention of education and is not presented as factual knowledge does not mean it can not teach the individual important abilities that could be decisive as far as survival is concerned.

---

<sup>1</sup> A more detailed examination of memetic value and its contributory factors appears in the article *The Attractive Error*, available at [alastairclarke.net/articles](http://alastairclarke.net/articles).

This means we also need to disinfect the erroneous *actions* of others to ensure we do not adopt their way of doing things. Consequently, an error on the page is treated no differently from a useless behavioural model from which we do not wish to learn, and either or both may constitute the misinformation the faculty addresses.

These few ideas are the essential principles of information normalization theory, according to which the system of humour is a remarkably compact adaptation, quite disproportionate in complexity to its contribution to human intellect. The equation  $h = m \times s$  represents the theory's universal rule, meaning there are no apologetic conditions for its application, and provided the reader is willing to examine the stimuli in question with an attention to detail then its use should prove straightforward for anyone with an elementary grasp of arithmetic and an average capacity for empathetic interpretation. To illustrate this, we will now turn what we have covered so far towards a common humorous event: a cartoon illustration caricaturing, for our purposes here, one of your friends, although the extent of familiarity is not important beyond the requirement that you possess a detailed knowledge of the subject's appearance. For the reward system of humour to function you need first to recognize an instance of misinformation, so if I show you a caricature of someone you have never seen before, you are less likely to be amused. A response may still be evoked based on a generic recognition of its distorted representation of a human being, but not its specific nature as a caricature. In our example the cartoonist has done rather well in your opinion, and you are amused by the picture. Your friend has a large hooked nose and this has been recognized and exaggerated in the portrait to draw attention to it, as would be expected in a successful example of this kind.

In order to relate this to the equation, we have to enter a value for the variable  $m$ , and to do this we need to decide to what misinformation the faculty is reacting. As tends to be the case with theories of humour, even simple explanations become a little trickier when the time comes to identify with precision the data that is being processed in each stimulus. One of the reasons for this is the undeniable subjectivity of the faculty, which reacts to protect the integrity of information as each individual perceives it. Consequently this means we can not now simply write out a list explaining what is or is not universally funny, tick the box next to humour and move on to another subject. Just as we do not all appreciate the same music despite hearing via the same mechanism, we do not all find the same events amusing, and nor do our own tastes remain static throughout time. Further,

two of us can laugh at the same data for quite different reasons, just as motivations for liking songs will differ between listeners. The reasons for this shifting subjectivity and intense individualization should be relatively clear. The freedom afforded by consciousness means our tendency to be amused is not only directed by our judgement of whether an obvious candidate should be registered as bad data or not; it is first and foremost affected by our selection of misinformation from the myriad streams of data that surround us. Having chosen points on which to alight, the biases and inclinations of our personalities will also then proceed to affect the memetic value we award them.

All of this means our illustrations are restricted to the theoretical, and should be considered models of interpretation for guidance only. Since the analysis of any instance of the faculty's activity requires intimate details of both the individual being amused and the circumstances of the misinformation's perception, employing a one-size-fits-all approach to humour can lead to some very misjudged conclusions. However, for the purposes of this illustration, we will presume a fairly limited and predictable reaction to our instance of pictorial caricature. Your friend's distinctive features have been exaggerated, and in this description lies all the elements we require to explain a typical example of the genre: first, the process of exaggeration is, by its nature, a misrepresentation. The norm, or the expected or accepted value, is increased. Your friend's large nose is there on paper, except that this is no accurately measured depiction. Instead, it is grossly enlarged, appearing as long as the rest of their head and shoulders put together. As students of humour it is necessary to overlook our own perceptions of memetic value, and it should therefore make no difference to our analysis whether we too would find this stimulus amusing. We may even find it offensive or immature, but as long as another individual has found it funny it is up to us to explain why, regardless of our personal responses.

Next we need to put a figure on this degree of exaggeration. Without wishing to get too bogged down in detail, before we can do so we briefly need to define the ranges of our variables. To keep the arithmetic as elementary as possible we will impose limits of 0 and 1 on the value  $m$ , and -1 and 1 on the variable  $s$ . The degree of misinformation will always be positive because we are only interested in stimuli in which it has been recognized, but the memetic value may prove negative if aspects of the event in which that misinformation arises register as particularly unappealing to the individual. In such a situation the misinformation is correctly identified

as such, but no disinfecting is required (and no humour is evoked) because the total value of  $h$  amounts to less than zero. If you find yourself the subject of an humiliating experience at which those around you laugh, you may not find it quite so amusing for this very reason.

Here, then, is your misinformation, taking the form of what you consider the significantly altered dimensions of your friend's nose, so we will award the variable  $m$  a medium-high rating of 0.6. Now we need to allocate to this particular aspect of misinformation the memetic value you associate with its presentation, allowing us to assess to what degree you are susceptible to being taken in by this infected data. To do so, consider again our description, in which we stated that your friend's *distinctive features* had been exaggerated. We have already decided that there is misinformation here, but are there any redeeming elements associated with or suggested by the picture that might lead you to take it seriously? In our example it appears so, because you think the caricaturist has done well, by which you mean that the *correct* features have been selected for exaggeration. Your friend does indeed have a large nose, and its presentation in an exaggerated form carries with it an implicit statement with which you agree. First-hand evidence tells you the depiction is half right, even though it is really quite unfair at the same time. Positively influenced by this commendatory metadata, your brain provides the illustration with a high memetic value of 0.75. Entering these figures into the equation produces:

$$h = (m = 0.6) \times (s = 0.75)$$

$$h = 0.45$$

We have our result, but how funny is 0.45? It is of course as funny as the two variables multiplied together will allow, which have both also been estimated here as representative impressions of how you feel as you view the stimulus event. Unless we undertake a careful and difficult calibration of our system according to the intensity of neurophysiological rewards, this abstract description of an instance of amusement will only provide an answer by comparison or via assessments based on personal estimation. The result is above zero, which means a response has been engendered, and is towards the middle levels of the internal pleasure associated with humour. With this unscientific, unmeasured but theoretically determined figure in mind, we now have a reference point from which to analyze difference.

Next consider an alternative attempt at the same image, in which the caricaturist has reduced the exaggeration of your friend's nasal features. As a consequence you recognize lower levels of distortion in this version, which you rate as 0.3. However, the pay-off arising from this alteration in the artwork is that it is now more convincing as a representative portrait, and the memetic value you identify within it increases. The caricaturist is still right that your friend's nose is a bit big, but you also award this second picture some merit according to realism, and the memetic value hits 0.85:

$$h = (m = 0.3) \times (s = 0.85)$$

$$h = 0.255$$

Even though this image is more useful as a source of information for those attempting to identify your friend, the two factors in combination produce a lower result than for the first version, and you find it less amusing as a consequence. Humour does not exist to recommend what is sound, but to encourage us to reject that which is not. We can continue this trend by asking the artist to sketch one final picture as a natural portrait. The elements of caricature are removed, and the nose assumes a dimension and profile similar to those we perceive in real life. Presuming the artist has done a good job there is now no noticeable misinformation, only memetic value, which has risen further. This is a good depiction, so we award it an impressive 0.92 for the variable  $s$ , but remain entirely unamused:

$$h = (m = 0) \times (s = 0.92)$$

$$h = 0$$

Of course, since there is no misinformation it makes no difference what memetic value we attach to the image as far as this particular faculty is concerned. No humour can arise, and the system remains inactive.

This is clearly a simplistic example. To begin with, we have not addressed how another individual might find your friend's nose amusing with or without the caricature. Secondly, there is only one bit of misinformation being assessed within the stimulus, whereas in practice multiple instances would often function in combination. Further, in our example the memetic value is determined by a single factor, and that one is consistently positive. Instead, there will nearly always be several, and each one can

be chosen from a countless and potentially unlimited range. Some, like novelty, may be relevant to the entire population, whereas others will be highly personal in nature. We have chosen representative accuracy as a memetic determinant in this example, but another individual may neither care about nor register its presence in their assessment, instead accepting that if the correct feature has been exaggerated and the intended subject is identifiable, the caricature is still sufficiently convincing to amuse them.

One consequence of the multiplicity of factors affecting memetic value in most stimuli is that the inverse correlation that was exhibited between the values of  $m$  and  $s$  in our example is not to be expected as the norm. With a range of positive and negative factors one may compensate for another, but the calculation will inevitably become a little more complicated than I have suggested when all the various influences over a person's attraction towards one unit of information are taken into account. If you like me, for example, you may award a greater memetic value to what I say regardless of its content, but, conversely, if you are bored by the subject matter of my conversation this may have the opposite effect. The range of possibilities for the manner in which misinformation is judged are only bound by the individual's perceptions.

If this is the case, considering the unlimited number of factors that could potentially contribute to the memetic value of a stimulus, how does the faculty know which ones to take into account? The answer is simple: it does not and does not need to. Our schematic analysis requires that we identify the major factors that generate memetic values because we need to record their effects in order to clarify why the stimulus is being found amusing, but all they really add up to is the overall impression the individual receives of the misinformation. This may be positive, negative or anywhere in between, but it will always be automatic to conscious perception. It is this reaction that is fed into the system, and there is no need for the faculty of humour to undertake an assessment of its own.

How, then, does this small equation slot into the remaining neurological apparatus we know as humour? I have claimed that  $h = m \times s$  describes the faculty's pivotal workings, yet it only appears to explain the process by which the intensity of the reward is calculated. It may therefore seem reasonable to assume that there must be associated engines performing other functions elsewhere, but the system has no need of them, and remains contained within this single activity. Humour is a reward mechanism, and if the conditions are met that evoke a response, the individual has accomplished some feat that has proved evolutionarily advantageous, and

all the necessary cogs for the determination of whether this is the case are covered by our universal equation. However, the process of undertaking the behaviour that is deemed to be positive for the species is not a part of the faculty's activity or even something that is supervised by it. It is the desire to experience endorphin-fuelled rewards that makes humour so effective, because to receive them it is necessary to see through misinformation to which we possess some degree of susceptibility. We do this by any means at our disposal, and the faculty of humour monitors our progress before stepping in to reward us for our successful efforts.

The benefits to the individual are immediately apparent. As we repeatedly engage these other means at our disposal, whether analysis, assessment, concentration or guile, we become more competent at their deployment. A defence system against deception and an enhanced ability to disinfect erroneous data are significant advantages to possess in a competitive environment, and it is not difficult to imagine their contribution to intellectual development. However, if information normalization theory is correct, on a much wider scale humour has stimulated a major advance more significant than the defence of any one individual, and it has done so by removing the brake on cultural expansion. Before the faculty evolved, the interference of error, deception and confusion in the chain of information dissemination would have slowed the rate of exchange significantly, and with faulty data an inherent problem of instruction the tendency to learn would have been retarded. The culture of the human race would have remained much poorer than it is today, repeatedly touching and receding from a ceiling through which it could not break. Considering the high susceptibility of memes to mutation, in the absence of a corrective mechanism the communication of behavioural instruction may well have remained restricted to a predominantly genetic medium.

Yet the system of humour proposed by this theory could have lifted such a cultural ceiling, allowing information to flow with a far lower incidence of memetic infection. More of it could have been passed accurately, and spontaneous mutations could have been eradicated when and where they arose. Better still, these improvements in error detection would have considerably reduced the time required to instruct others, and, once taught, the erstwhile student would have been able to preserve the integrity of their instructions and assume the role of teacher. With the simple process of amusement becoming associated with the recognition of beguiling misinformation, the rate of all cultural exchange would have increased markedly, and the complexities of the modern world would have drawn a significant step closer.

THIS ARTICLE PUBLISHED BY DOT NET PRESS  
COPYRIGHT ALASTAIR CLARKE 2010

WITH THE PERMISSION OF  
PYRRHIC HOUSE


DOT NET  
■ PRESS ■

[WWW.PYRRHICHOUSE.CO.UK](http://WWW.PYRRHICHOUSE.CO.UK)

[WWW.ALASTAIRCLARKE.NET](http://WWW.ALASTAIRCLARKE.NET)

TO CONTACT THE AUTHOR E-MAIL  
[EVOLUTION@ALASTAIRCLARKE.NET](mailto:EVOLUTION@ALASTAIRCLARKE.NET)





The full-length publication *The Attractive Error* is due for release by Pyrrhic House in 2012, and will be available in hardback and e-book versions. To pre-order a copy or for further information visit: [www.pyrrhichouse.co.uk](http://www.pyrrhichouse.co.uk).