

Double-Edged Words

Surfing the Web for the Brave New Words of Tomorrow

Words are consumer goods. Some are durable goods, unsexy but reliable, that last us a lifetime. Others are fashion accessories that come and go with the seasons, popularized by celebrities and intimately tied to the latest short-lived trends. Like the latest must-have products which marry innovation with tradition, new words are rarely made from whole cloth, but bring together existing words and ideas in a catchy streamlined form. Nowhere is this continuous process of verbal repackaging more evident than on the web, a linguistic bazaar where users eagerly peddle their latest homemade concoctions. In this chapter we examine how new terms – creative neologisms – can be coined as combinations of meaningful fragments torn from other, existing words.

Principles of Ergonomic Word Design

Words are everyday things, as central to our daily lives as the clothes we wear, the tools we use and the vehicles we drive. Words are everywhere, and just like the consumer goods that stuff our closets yet whet our appetites to buy even more, we never seem to have enough words. Indeed, like consumer goods, words go through fashion cycles, becoming hot and painfully trendy one moment, when it seems that the same “buzz” words are on the lips of everyone who counts, only to be forgotten at the back of our lexical closets the next, waiting to be recycled someday as something retro and cool.

As man-made objects, words and phrases are subject to many of the same design principles as the consumer artefacts that compete for our attention in the marketplace. In his book *The Psychology of Everyday Things* (later reissued as *The Design of Everyday Things*), Donald A. Norman identifies two key principles of artifact design: visibility and mapping. A good design makes it easy for a user to mentally visualize, or conceptualize, the inner workings of a product, while a bad design causes a user to construct an inaccurate conceptual model that leads to misuse of the product and inevitable human

error. If well-designed, the external elements of a product will yield a natural mapping to its internal functions, but if badly designed, the mapping between appearance and function will be confusing and counter-intuitive. Quite simply, the buttons and switches will not perform the functions a user believes they should. These principles are just as applicable to novel linguistic designs – such as neologisms, witticisms and revitalized clichés – as they are to refrigerators or car stereos, since it is also desirable that the products of linguistic creativity should exhibit not only visibility of meaning, but a natural mapping of linguistic form to conceptual structure.

Manufacturers place new kinds of ovens, televisions and automobiles on the market all the time, but users do not need to relearn basic behaviours like baking, watching TV or driving to work. These new products are usually variants of existing models, adding new functionality and subtlety to familiar forms that retain their underlying structures. Likewise in language, new coinages frequently borrow the form of existing words and phrases, allowing a user to reuse the same underlying conceptual model. When presented with a novel coinage like “*Ghost airport*”, we don’t try to build a new conceptual model from first principles; rather, we reuse the conceptual model of “*Ghost town*”, by accepting that an airport is sufficiently similar to a town for the meaning of “*Ghost airport*” and “*Ghost town*” to be analogous (towns and airports tend to be filled with people, thoroughfares and businesses, while “ghost” variants are empty and desolate). Likewise, you may have never encountered the term “*Twitchhiking*”, but the word shares enough structure with “hitchhiking” to strongly suggest that the conceptual model for the latter can safely be reused. Depending on your technological savvy, you might guess that “Twitch” is a blend of “hitch” and “Twitter”, and integrate your knowledge of this new form of communication into the conventional model of hitchhiking. These new words really are *blends* in the sense explored by Gilles Fauconnier and Mark Turner, both on a formal level (that is, on the level of word forms, as in “*Twit*” + “*chhiking*”) and on a conceptual level (since aspects of the underlying ideas are blended as well).

Variation of an established convention is a common strategy in linguistic creativity, but not every variation is creative. For instance, variations in how a word is spelled or pronounced can yield a more or less creative pun, but random typing errors are highly

unlikely to yield anything we might consider creative. We should allow for serendipitous creativity that is unintentional, or the product of purely random combination or mutation, but almost all random variations will be uncreative, or else the very idea of creativity becomes devalued. What gives a linguistic variation its creative value is the transformation it yields in our understanding of the underlying idea. Suppose, as the designated driver at a party, you ask for a soft-drink and a friend offers you a “*Virgin Mary*”. Knowing that a “*Bloody Mary*” is a cocktail made from tomato juice, vodka and Worcester sauce, you will likely assume that a “*Virgin Mary*” is a soft variant of this otherwise alcoholic drink. Though a creative and economic use of language, the blend “*Virgin Mary*” has one foot firmly planted in the past, in the form of established linguistic convention, while the other foot leads this stock phrase in new and creative directions. Who knows what “Mary” refers to in either version of the cocktail? What matters is that this word provides a convenient bridge between the familiar and the novel. Interestingly, though it is the word “Bloody” in the conventionally alcoholic version of the cocktail that makes way for the word “Virgin” in the new, you are unlikely to assume that it is the blood-red tomato juice that is replaced, but the alcoholic vodka. In this case, “Virgin” suggests “chastity”, which suggests “abstinence”, which suggests “temperance”, which implies a lack of alcohol. This linguistic blend works about as well as the underlying cocktail, provided that we do not pay too much attention to the contorted mapping between surface form and conceptual structure.

In contrast, no such contortion is needed in the Australian name for this non-alcoholic cocktail, a “*Bloody Shame*”. This second variation is arguably the more ingenious and humorous of the pair, for a number of reasons: first, because “bloody shame” is already a familiar phrase in English, and so this variation establishes a punning relationship between the new cocktail and its source norm; second, because this pre-existing phrase has a negative connotation, of “regret” or “tragedy”, and this allows the variant name to express a negative view of the underlying concoction; and third, this negative perspective also expresses a strong cultural preference for alcohol that serves to reinforce the stereotype of the hard-drinking Australian. Drinkers who order a “*Bloody Shame*” thus communicate a disdain for their own choice while implying a desire to order something else – something a good deal more alcoholic – and order it in a way that

humorously seems to crave our sympathy. Indeed, because the phrase “*Bloody Shame*” has connotations of tragedy, it also works as a form of epic irony when used as the name of an alcohol-free cocktail. To sum up, “*Virgin Mary*” and “*Bloody Shame*” are both variants on a familiar phrase that use a single-word replacement strategy to shoehorn even more meaning into the same linguistic structure. But all variants are *not* creatively equal: while both variations work quite well, the latter achieves the greatest degree of creative duality, compressing multiple levels of meaning and perspective into a simple two-word name. Viewed as carefully manufactured products of human ingenuity, a “*Bloody Shame*” clearly exhibits the better design of the two names.

Designer Words

New words can be formed any way you like. They don’t even have to be pronounceable, nor for that matter do they have to be made from real letters. Remember the pop-star *Prince* and the squiggle he chose as a new name to exasperate music journalists everywhere? Yet there is a considerable distance between the ideal of inventing a fancy new word and the difficult reality of persuading everyone else to use it so that it becomes linguistic currency. Prince may have wanted his new stage name to grab the public’s imagination, but journalists and fans alike preferred the alternative (and somewhat self-defeating) moniker *The Artist Formerly Known as Prince*. The lexical gloss may have been longer, but it was pronounceable and said a lot more about its referent than any bizarre doodle ever could. In other words, the long-winded alternative exhibited Norman’s ideas of visibility and mapping.

If your new word is intended to denote an idea that is best described as a combination or blend of other ideas, then you could do a lot worse than using a blended word that combines lexical elements associated with those other ideas. The mathematician and wordsmith Lewis Carroll assigned a memorable label to these combination words, though sadly not one that followed the same principle of word-level blending. Carroll called these words *portmanteau* words, named for the double-pocketed satchels that allow their users to organize different facets of their life into separate compartments. As he describes them in *Through the Looking Glass*, a portmanteau has

“two meanings packed up into one word”. The usefulness of a new portmanteau is generally a function of the utility of the idea it describes, and of the importance of communicating this idea (and all that it entails) concisely and without ambiguity. When the idea is a category, the usefulness of the portmanteau will generally be a function of the size of the category, and of the importance of not confusing this category and its members with those of a rival category. The word “*bromance*”, for instance, describes a genre of comedic movies of the boy-boy (rather than boy-girl) variety, in which two male characters share a close friendship that borders on platonic romance. The increasing popularity of this genre, and its strong similarities (but obvious differences) with the romantic comedy, have contributed to the popularity of the “bromance” label amongst critics and viewers alike. While most viewers like to be pleasantly surprised when they go to the movies, most also prefer movies that can be accurately summed up in a word or two. Popular creativity must work with the familiar to make the novel acceptable. As Hollywood mogul Louis B. Mayer once put it, “Let’s have some new clichés”.

The portmanteau principle is responsible for some truly horrible words when used lazily, though this is largely a matter of personal taste. For instance, the word “brunch” seems fine, since the “unch” of “lunch” is pleasingly suggestive of the word “crunch”, but “spork” seems an ugly word, perhaps because we rarely use disposable plastic spoon-forks to eat anything that is actually worth eating. The word “sharpedo”, denoting a torpedo for killing sharks, was either coined by a six-year old or by a mental defective. The triple-threat portmanteau “turducken”, denoting a turkey stuffed with a duck that is in turn stuffed with a chicken, violates an unspoken culinary principle that every commercially-savvy chef should know: the time course from food to excrement is all too short as it is, and the word “turd” has no place at all on a menu. Yet if used properly, with insight and respect for the words involved, a portmanteau word can combine high visibility with a natural mapping of surface elements to underlying ideas, making portmanteaus the perfect designer words. The word “affluenza”, for instance, concisely captures the idea that the affluence of others is both contagious (“keeping up with the Joneses”) and personally destructive, like a harmful disease. Likewise, the word “malware” does more than save a speaker a syllable or two; it gives an evocative name to a whole category of software nasties – from viruses to Trojan horses and worms to

adware, spyware and crippleware – that can malevolently inflict damage on a computer. The prefix “mal” nicely suggests that bad juju is afoot, while the affix “ware” does an excellent job of designating “malware” as a category with a broad membership, on roughly the same level as “hardware” and “software”.

Lewis Carroll delighted in creating apparently nonsensical words using the portmanteau principle, such as “slithy” (“lithe” and “slimy”) and “snark” (“snake” + “shark”), which are lexically suggestive if not exactly semantically transparent. Modern uses of the portmanteau principle generally aim for greater transparency, allowing a reader to infer the constituent words (and thus, ideas) from which the neologism is blended. The historian Niall Ferguson has coined the portmanteau “Chimerica” to describe the heavily inter-dependent relationship between the U.S.A. and the People’s Republic of China. As Ferguson puts it, America and China are no longer two distinct countries from an economic perspective, but one blended economic whole deserving of its own name, *Chimerica*. Ferguson’s coinage results in a rather ugly word, but it is a word with some interesting properties nonetheless. For one, the word *Chimerica* resembles the Greek word “Chimera”, a mythical monster that combines parts of other fabulous beasts, such as the body of lioness and a tail with a snake’s head. The word “Chimera” is also used in modern genetics to describe a single organism with genetically distinct cells from two different zygotes. This is essentially what a portmanteau word is: a neologism that results from the cross-breeding of words. Another interesting property, then, is the suggestion of conceptual unity that arises from the structure of a portmanteau word: the tight lexical integration of two distinct word-forms into a unified lexical whole suggests an equally tight integration of ideas at the conceptual level.

Speakers do not have to like a portmanteau to find it useful and strangely compelling. Here is the *Time Out* guide to Mumbai and Goa discussing that much-used portmanteau of “Bombay” and “Hollywood”, “Bollywood”:

“The term Bollywood is despised by many in the Indian film industry, not least for defining Indian films in relation to Hollywood, but it remains an unrivalled catch-all phrase for describing the farrago of emotion, action, song, dance and humour that animate almost every Hindi film.”

Humorous effects can arise when integration at the lexical level forces together ideas that one might consider incompatible at the conceptual and pragmatic levels. Consider the word “*Feminazi*”, a portmanteau coined by the political scientist Tom Hazlett and popularized thereafter (to controversial effect) by Hazlett’s friend, conservative radio host Rush Limbaugh; clearly one doesn’t have to be a genius to coin a catchy portmanteau. You might imagine that “*Feminazi*” was coined to equate the strident expression of feminism with the Socialist Nationalism of Nazi Germany, but Limbaugh and friends have used it to characterize almost any woman with progressive views. Though used for dubious social purposes, the word is nonetheless inventive and not without humour, which arises here from the clash of semantic frames associated with the terms “Feminist” and “Nazi”. These frames can be reconciled, at least superficially, by seeing each type of agent as a zealous advocate of a social philosophy. The word is cheeky and childish and a little fun even if it is extreme (or perhaps *because* it is so extreme). Portmanteau words have the power to unite incongruous (but *appropriately* incongruous) ideas in a potent lexical form, and Hazlett and Limbaugh’s “Feminazi” is the lexical equivalent of daubing a swastika on Gloria Steinem’s front door.

Bubbling Under, Boiling Over

Limbaugh occupies a position that most aspiring wordsmiths will surely envy: his radio show, though most certainly an acquired taste, is highly influential, allowing him to hammer home any new word that takes his fancy. Of course, “Feminazi” was first coined to bait feminists and people with progressive views everywhere, and these are precisely the kind of people that do not listen to Limbaugh’s show. Nonetheless, his show provides a loud and persistent voice in the political discourse of the day, and new coinages that are sufficiently catchy can quickly make it into the linguistic mainstream. Critics who responded to the irony of comparing feminists to Nazis (since the latter systematically stripped people of their rights, while the former fought for these rights) merely succeeded in giving the term greater air-play and an even wider reach, well beyond the confines of the conservative audience it was designed to tickle. The new word, designed to entertain some and anger others, did not so much reach a tipping point as a boiling point, at which

stage Limbaugh pulled back on its usage.

Other new coinages experience exhibit quite a different trajectory as they make their way into popular usage. The portmanteau “*metrosexual*” was coined by journalist Mark Simpson in 1994, in an article in the British newspaper *The Independent*. Simpson defined a metrosexual as a “single young man with a high disposable income, living or working in the city”, noting that for these young city dwellers, access to trendy shops, cool bars and male vanity products is a more defining feature than sexual orientation: the key distinction, Simpson argues, is not same-sex (homo) or different-sex (hetero) but city (metro) versus non-city. His timing may have seemed apt, since the 1990s witnessed a rise in “lad culture” in Britain, in which well-heeled and well-groomed young men were targeted by almost as many lifestyle magazines as young women. Nonetheless, the term failed to take off in a significant way, perhaps because the kind of individual it was designed to describe had not yet become an identifiable urban stereotype. Undeterred, Simpson re-launched the term in 2002, in an online article for the website *Salon.com*. By this time, one could name a variety of high-profile role models for the term, from the actor Jude Law to the footballer David Beckham. Simpson cited the latter as the very prototype of a metrosexual, prompted no doubt by Beckham’s love of clothes, perfume and jewellery, and by his hawking of various grooming products on international TV. Second time out, Simpson’s portmanteau hitched itself to a cultural phenomenon that made it a big hit, spawning a variety of derived terms, from the obvious “*metrosexuality*” to the less obvious “*übersexual*” and “*retrosexual*”. The latter two were seemingly spawned as a cultural reaction to the prevalence of metrosexual coverage in the media, and as the prefix “retro” implies, “*retrosexual*” denotes the masculine ideal before the rise of metrosexuality, as typified by actors with a low preen-factor like George Clooney.

“Seemingly” is the key word here, since much of the so-called debate about metrosexuality was conducted in newspapers and fuelled by writers hoping to surf the Zeitgeist with a new book about hot social trends. In truth, newspapers like nothing better than a convenient and reductive label for that nexus of people or behaviours that are currently in or out of favour. New words make for very convenient labels, and even portmanteaus – which are constructed from bits of existing and well-worn words – hint at deep insights (a new category of X has been discovered!) and give the suggestion that

more is being communicated than is actually the case. This deliberate under-specification allows a new word to act as a convenient handle on a cloud of related observations and unstated prejudices, with greater headline-friendly ambiguity than a descriptive phrase could ever muster. New words begin life with one meaning but can grow into another.

The suffix “sexual” is a promiscuous partner in the creation of many novel portmanteau words. The most frequent non-standard pairings include “hobosexual” (denoting the opposite of a “metrosexual”: men who take little obvious pride in their scruffy appearance), “autosexual” (not a lover of cars, but one who prefers self-gratification), “biosexual” (a passionate nature lover), “electrosexual” (a nerd who prefers electronica, such as video games, to sexual gratification), “technosexual” (someone with a near-sexual lust for technology), “emosexual” (a moody, androgenous lover of emotional “emo” music), “petrosexual” (a “petrolhead”, a person obsessed with cars), “pomosexual” (a postmodernist who considers rigid sexual categories to be too structuralist for words; perhaps also a postmodernist with affectations that suggest the negative “homosexual” stereotype; flamboyant instances are sometimes called “a flaming pomo”), “prosexual” (someone who likes all forms of sexual gratification, and a believer in sexual intercourse for its own sake), “pseudosexual” (someone who disguises their real sexuality with deliberate miscues), “psychosexual” (which as an adjective refers to the psychological/Freudian aspects of sexual development; but which, as a noun, can refer to a person who derives a near-sexual pleasure from playing “mind games”, or more coarsely yet aptly, “mind fucks”), “robosexual” (one with an unhealthy fixation on robots), “slomosexual” (someone who enjoys slow sex, that is, sex in “slow motion” or “slo-mo”), and “zoosexual” (one with an unnatural love of animals).

When *CNN Money* announced the arrival of “*jetrosexual*”, a new variant on “*metrosexual*” coined by the airline *Virgin Atlantic*, with the headline “There’s a new type of ‘sexual’ in town”, it was clear that the game was up for the “sexual” suffix. No longer a productive affix for coining playful new words, it had become the morphological equivalent of a snowclone. According to a *Virgin* press release, a *jetrosexual* is a traveller who has a favourite airport, can order a beer in six different languages, and who may occasionally travel in economy class, but only as a self-imposed lesson in humility. We

must assume that a jetrosexual is also a special kind of metrosexual, for how else are we to understand the superfluous “ro” between “jet” and “sexual”? This is wordcraft at its laziest, and violates both of Norman’s design principles of naturalness and mapping. The “ro” fragment is just a left-over, the lexical equivalent of a surgical sponge accidentally left inside a patient by a careless surgeon. The purpose served by each lexical fragment in a portmanteau may not be immediately obvious, but each *should* serve a purpose.

Two-Fisted Wordsmiths

Words that are constructed using the portmanteau principle can land a double-punch on their audience, bringing together two unrelated ideas or domains of experience in a striking new form. Yet, as we can see from the previous examples, the portmanteau *principle* is not really a principle at all, but a sliding scale of compositionality and reuse. Some portmanteaus deserve to be considered pure examples of the form, combining clippings of other words that cannot stand alone but which work together in harmony as part of a novel word-form. For instance, the words “brunch”, “spork” and “affluence” are all made from elements that have no meaning in their own right, but which serve to evoke other words and other meanings in context. Words like “metrosexual”, “retrosexual” and “technosexual” muddy the waters somewhat; as we have seen, even CNN now uses “sexual” as a stand-alone noun, one that denotes a personality-type with a catchy portmanteau label, while the prefixes “metro”, “techno” and “retro” are all used in English as stand-alone nouns and adjectives. Many portmanteaus are semi-pure, like “Bridezilla”, combining a stand-alone word (“Bride”) with a suggestive affix (“-zilla”, implying “monster” via “Godzilla”) that is itself incomplete as a word. The widespread duality of existing words is one reason why purists adopt a strict clipping-based definition of portmanteaus. If “malware” is considered a portmanteau, then why not “malfunction”? If “metrosexual” is considered a portmanteau, then why not “metroplex”?

In truth, many new words – and perhaps most *successful* new words – are marriages of multiple influences. When a new word is coined to replace a longer phrase, we can expect the new coinage to incorporate influences from the component words of the original phrasing. This is what we should expect if words are “soft” technologies that

obey Don Norman's intuitions about good design. Likewise, a catchy new word strives to satisfy dual competing goals: to say something obvious about the idea it denotes, so that an audience can immediately see its suitability as a label for that idea; and to say something non-obvious about this idea, so that the word is interesting enough in its own right to be used and passed along by this audience. As with new technologies, speakers are avid consumers of new words, and can get a thrill from being an early adopter of the latest vocabulary, of being the first in their peer group to use a clever coinage.

Take the word "astronaut": when America found itself lagging behind the Soviets in the race for space, the newly formed NASA cast about for an imagination-grabbing name for the intrepid pilots who would spearhead America's manned space program. NASA did not coin the word "astronaut" – it had been used in fictional contexts some years earlier – but it did popularize the word and burnish its heroic appeal. New terms were required for new times and new endeavours, and the popularity of "astronaut" grew in tandem with the public's appetite for all things technological and futuristic. The word seemed shiny and new, but it was not cut from whole cloth; rather, it had an obvious precedent in "argonaut", which described the heroes of the Argo who followed Jason in search of the golden fleece, and in "aeronaut", which was first coined to describe early ballooning pioneers. The "naut" of "argonaut" is literally apt, denoting a sailor aboard a ship (it is the same "naut" of "nautical"), while it is metaphorically apt for astronauts and aeronauts. But the "astro" of "astronaut" is fantastical and grandiloquent, meaning "star", though the "cosmo" of "cosmonaut" (meaning "universe") is just as over-reaching. Words like "astronaut" and "cosmonaut" are not portmanteaus, but they are clearly the product of lexical integration – the fusing together of separate word elements to create a new and integral whole.

A language can be seen as a rich set of Lego bricks. Bricks of different colours, shapes and sizes can be combined to create elaborate structures, just as the words or morphemes of a language can be combined to create elaborate phrases, sentences and texts. Some bricks sit so well together in so many different contexts that we may begin to think of them as composite super-bricks, and over time, we may even forget their composite nature. To do something creative and surprising with our brick set, we can

look to break up these super-bricks into their component parts, and recombine these parts in pairings that have not been used before, in the hope of creating new super-bricks that are just as popular and useful. These new super-bricks can be constructed in two different ways. In the first approach, one of the components in an existing super-brick is replaced with another brick with much the same shape but with enough additional properties to motivate the replacement. Take the word “*modelizer*”, which Candace Bushnell introduced in 1995, in an article in *The New York Observer* (the article was later re-used, as a text-level super-brick, in Bushnell’s novel *Sex And The City*):

“Modelizers are a particular breed. They’re a step beyond womanizers, who will sleep with just about anything in a skirt. Modelizers inhabit a sort of parallel universe, with its own planets (Nobu, Bowery Bar, Tabac, Flowers, Tunnel, Expo, Metropolis) and satellites (the various apartments, many near Union Square, which the big modeling agencies rent for the models) and goddesses (Linda, Naomi, Christy, Elle, Bridget).”

The word became popular enough to earn its place in the Oxford English Dictionary in 2003. Since “modelizer” can be seen as a contraction of “model womanizer”, one might think of it as an impure portmanteau of sorts, though the clipping “izer” is not especially suggestive of “womanizer”. It is more insightful to see “modelizer” as a *specialization-by-substitution* of “womanizer”, in which the brick “woman” has been replaced with the specialization “model” (following the stereotype that models are usually women, not men). This process of specialization-by-substitution can be cumulative, and Bushnell goes on (in the novel) to speak of “supermodelizers”, men who are so adept at snaring beautiful women that they exclusively focus on supermodels.

In the second approach, super-bricks are broken down into their component parts, to yield an inventory of reusable elements that can then be re-combined in many different ways. From this perspective, words like “astronaut”, “heterosexual” and “metropolitan” are the ideal super-bricks: easy to identify, and easy to break into reusable parts. This ease is due in large part to the meaningfulness of the components, for though fragments like “hetero”, “cosmo”, “politan” and “naut” do not typically stand alone as words, they

communicate specific meanings on their own terms, rather than by allusion to the larger words in which they can appear. As we have noted, “cosmo” signifies “universe” in Greek, “astro” signifies “star”, “naut” signifies “sailor” or “voyager”, while somewhat surprisingly, “metro” is Greek not for “city” (that distinction goes to “polis”, from which “politan” is derived) but for “mother” (hence a “metropolis” is the “mother city”, or as well-known dictator, now deceased, might put it, “the mother of all cities”). These Greek-derived lexical components are highly-combinatorial Lego bricks that play very well with others, allowing a wordsmith to fuse them into many novel arrangements.

Lego BrainStorms

Wordsmiths typically coin a new word to signify a specific idea in a particular context. As such, the process of word creation is usually tailored to the task at hand, making it a most difficult process to accurately model and simulate. For if we cannot recreate the ephemeral contexts that motivate the coining of new words, we cannot meaningfully recreate the mental processes of creative individuals as they react to the demands of these contexts. After all, who can say what does through the mind of such an individual when faced with a lexical choice that is challenging enough to necessitate a completely new word? But if individual, once-off creativity in an ephemeral context is too under-determined to make for a compelling computational model, we can look instead to a slightly more artificial yet creatively meaningful task: that of brainstorming a wide variety of new words for an equally wide variety of possible meanings.

Picture the scene: a group of more-or-less creative individuals are spit-balling in the conference room, trying to suggest as many meaningful but creative new words as possible. A fun exercise, perhaps, but why would one ever pay expensive consultants to play word games? Surprisingly, this kind of wide-ranging brainstorming session is more common than one might think, and it is not unusual for a large company to strive to create and maintain an inventory of catchy product names in advance of any new products that might actually bear those names. Certainly, the marketing wizards at Greek airline *Olympic* could have benefited from the advice of a professional name consultant when it decided to brand its frequent flyer program with the ill-conceived name *Icarus*

(the Greek aeronaut famous for falling into the sea after his wings fell apart, so hardly a name to inspire confidence in travellers). For instance, the pharmaceutical giant Pfizer maintains an inventory of candidate names for the new drugs that may eventually pop out of its development pipeline. These names are first brainstormed without prior knowledge of the specific drugs that they may one day grace, and ranked according to a range of product-independent criteria, such as “is the name novel?”, “is it catchy?”, “does it suggest positive qualities?” and “does it mean anything offensive, taboo or silly in any of the languages/countries in which it may be marketed?”. Thus, while one might think that Pfizer coined the brand name “Viagra” for its best-selling impotence drug because this name is especially resonant of the idea of male sexual “vigor”, or because it suggests tremendous natural power due to its similarities with the word “Niagara”, the name was actually coined long before Pfizer ever developed the drug or identified its value in treating sexual dysfunction. Though brainstormed *before* product development, the name was chosen from a list of candidates *after* product development, which is why it seems such an especially apt label. If chance favours the prepared mind, the pre-generation of a wide-range of well-formed candidates can yield real benefits in any creative choice.

As an experiment in word-level creativity, let’s simulate the brainstorming of new words by combining meaning-laden fragments of existing words. This will be more general than the brainstorming of product-names, but will follow much the same process: new words will be generated only when they are suggestive of sensible meanings that could plausibly benefit from a convenient single-word label. These fragments are “meaning-laden” in the sense that each fragment – either a word prefix or a suffix – has its own established meaning; thus, we shall use fragments like “astro” to mean STAR, or “techno” to mean TECHNOLOGY, or “naut” to mean TRAVELLER, rather than the arbitrary “sp” for SPOON, (as in “spork”) or “unch” for LUNCH (as in “brunch”). This idea of using word or name fragments to coin creative new labels is a well-established one. Computer scientist Wlodzislaw Duch refers to these fragments as *wordels* or *word elements*, by analogy with *pixels* and *picture elements*. The notion even forms the basis of a commercial software package called *NameRazor* for helping people coin their own product and website names. *NameRazor* refers to its inventory of mini-names as *namelets*, and has the curious distinction of having suggested its own name, with a little

help from its creators.

To set up our brainstorming experiment, we need to resolve three practical issues:

- a) In which existing words will we find the most reusable word fragments?
- b) What is the best way of splitting these existing words into their reusable parts?
- c) How do we assign meaning to these reusable parts?

The second issue can be resolved by assuming that the most natural place to break a word in two, for purposes of generating a reusable prefix and suffix, is also the most natural place to break a word in two for purposes of hyphenation. Consider how people break up a word when it is too big to fit at the end of a line of text: writers do not insert the hyphen just anywhere, but draw upon complex intuitions about spelling, pronunciation, typography and lexical construction to place it between the most meaningful sub-parts of a word, such as after a common prefix or before a common suffix. Not everyone hyphenates sensibly, of course, but looking through a large text corpus we can expect the most frequent hyphenations to be the most sensible. In such a corpus, we are far more likely to encounter “God-zilla” than either “Godz-illa” or “Go-dzilla”, while “astro-naut” is much more likely than either “astron-aut” and “astr-onaut”. Since the most frequent is the most widely favoured and thus the most sensible, we need not articulate the specific intuitions that writers use when breaking up words: we can always learn specific hyphenation preferences for specific words from a large text corpus like the web.

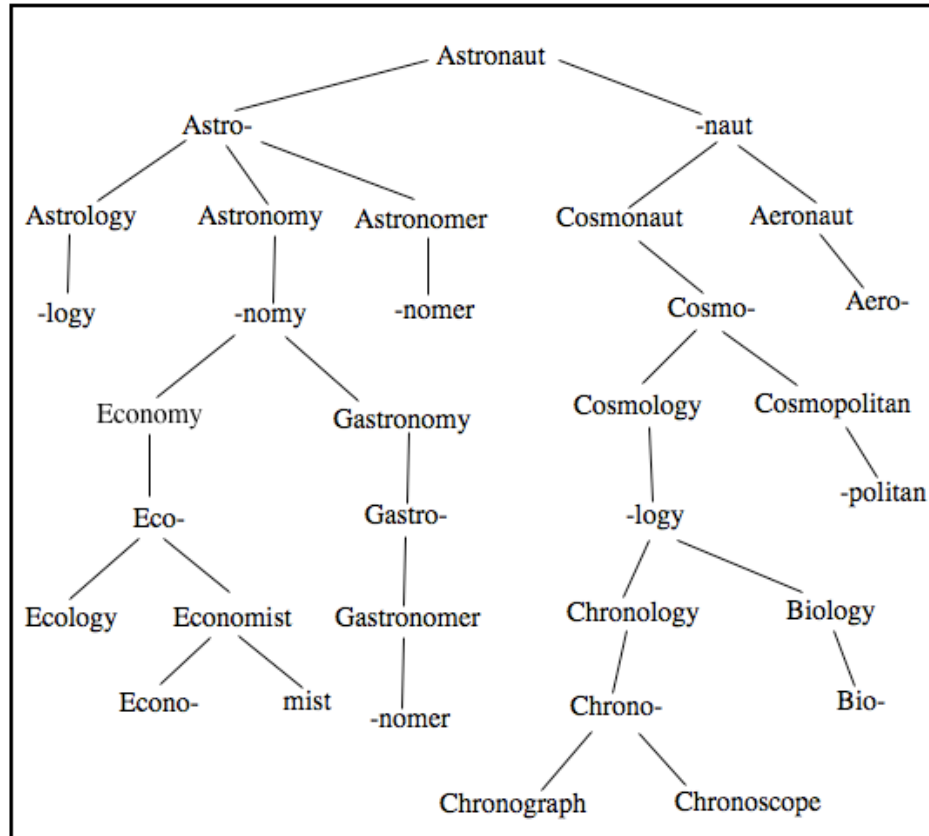


Figure 1. *Starting from a single seed, one can reach a wide range of other similarly-structured words with a shared prefix or suffix. This allows us to identify a large inventory of reusable prefixes and suffixes.*

The first issue – which words to break up into reusable fragments – now becomes a matter of personal preference. Suppose we favour words like “astronaut” that combine two well-established fragments of Greek origin. Using a small set of exemplars as a breeding stock, we can find other words of a similar structure and origin, by first dividing the exemplars into their separate prefix and suffix components, and by then recursively seeking out other words with a shared prefix or suffix. Figure 1 shows this *divide-and-recruit* process at work, starting with a single exemplar “astronaut”. The suffix “-naut” leads to the words “cosmonaut” and “aquanaut”, which yield the new prefixes “cosmo” and “aqua”; “cosmo-“ in turn leads to “cosmopolitan” and “cosmology”, which yielding the new suffixes “-politan” (which later provides “metro-“) and “-ology” (which is a mother-lode of Greek prefixes); and so on. Figure 1 shows that the same suffixes and

prefixes are likely to be found over and over again, such as “-logy” (reachable via “astronaut”→”astro-“→”astrology”→”-logy” and “astronaut”→“-naut”→”cosmonaut”→”cosmo-“→”cosmology”→“-logy”) and “-nomer” (via “astronaut”→”astro-“→”astronomer”→“-nomer” and via “astronaut”→”astro-“→”astronomy”→“-nomy”→”gastronomy”→”gastro-“→”gastronomer”→“-nomer”). As one might expect, the more frequently a particular fragment is used in different words, then the more often we are going to encounter it during this *divide-and-recruit* process. But this redundancy does not lead to wasted effort, since the number of times we find the same suffixes and prefixes is a good indicator of how useful those fragments will be in the creation of new words.

This reuse means that if we are willing to search deep enough, then even a small breeding stock of words can yield a relatively large inventory of useful word fragments. For instance, if we start with the word “psychology”, which gives us up front the widely used fragments “psycho” and “logy”, a search that is 10 levels deep will yield an inventory of over 900 fragments, such as the much less frequent (but no less meaningful) “caco-” and “-lyte”. As we might expect, different words from the same language family exhibit a strong overlap in the fragments they retrieve. The Greek-derived fragment “-phobia” allows us to identify 378 other fragments (when searching 5 levels deep), 56% of which are also retrieved if we start from “-logy”; the fragment “-scope” allows us to identify 333 other fragments, 95% of which are retrieved by “-logy” or “-phobia”; and the fragment “-meter” leads to the identification of 212 other fragments, all of which are reachable from either “-logy”, “-phobia” or “-scope”. Once again, this tendency for overlap is a sign of linguistic coherence, which suggests that these word fragments should interact well to frequently produce meaningful, if novel, combinations.

With a large frequency-ranked inventory of word fragments in hand, we can now consider how to assign a conventional meaning to each, such as STAR for “astro-” and FEAR for “phobia”. We can imagine a variety of tricks for doing this, such as using an online etymological dictionary, but the quickest and most reliable approach is also the most old-fashioned: for quality outputs we need quality inputs, and there is simply no substitute for the manually assigned meanings of a native speaker who actually knows what he or she is doing. For instance, we can hand-craft the mappings in Figure 2:

-logy:	<i>study, discipline, system</i>	chrono-:	<i>time</i>
-scope:	<i>display, monitor, viewer</i>	gastro-:	<i>food, cooking, eating</i>
-naut:	<i>traveller, sailor</i>	necro-:	<i>dead, death</i>
-polis:	<i>city</i>	pyro-:	<i>heat, fire, flame</i>
-metry:	<i>measurement</i>	psycho-:	<i>mind, mental, crazy</i>
-mancy:	<i>magic</i>	geo-:	<i>place, location</i>
-pyle:	<i>opening, gap, vent</i>	cryo-:	<i>cold, ice</i>
-nym:	<i>name</i>	hydro-:	<i>water</i>
-trope:	<i>seeker, follower</i>	helio-:	<i>sun</i>
-glyph:	<i>marking, symbol, icon, rune</i>	phono-:	<i>sound, hearing</i>

Figure 2. *Different word fragments, to which coherent meanings are assigned.*

We are dealing with a simple inventory of about 400 fragments here, so we may as well do it properly, and manually assign the meanings that are most conventionally associated with each word fragment. For instance, we can just skip over the strict dictionary meaning of “gastro-“ (denoting BELLY or STOMACH, as in “gastroenteritis”) and use the less biologically-oriented glosses FOOD and COOKING instead. Likewise, we can overlook the strict dictionary definition of the prefix “metro-“ (denoting WOMB or UTERUS, as in “metrorrhagia”, and thereby denoting MOTHER in “metropolis”), to assign instead the more conventional (if strictly incorrect) modern associations of URBAN and CITY (which makes “metropolis” something of a redundant combination). Our goal in this brainstorming experiment is to coin novel words that might possibly achieve common currency, rather than technical jargon for the classically-minded few. While many will undoubtedly be of the latter cast, we can nonetheless hope that some will be of the “metrosexual” variety; fortunately, since our fragments come from existing words, ill-begotten mutants such as “jetrosexual” are off the table.

Let the Brainstem-storming Begin

The Design Conspiracy is a trendy London-based consultancy firm with strong opinions about what constitutes good – and bad – design. In 2003, the company grabbed headlines when a spoof web-site offering contrived company names on demand was taken seriously by some businesspeople. The site, a publicity stunt named *WhatBrandAreYou.com*, works a little like those tell-your-fortune machines that used to be popular in amusement arcades and funfairs: the user is asked for certain personal details, such as whether you or your company can be described as “passionate”, “dynamic” or “client-focused”, but these are then promptly ignored. Like those arcade machines, this spoof site works with a stock of canned outputs (about 150 in all), from which the site chooses randomly for each user. Its outputs include credibility-stretching names like *Ovisovis*, *Amplifico*, *Bivium*, *Libero*, *Vulgo*, and *Ualeo*, each produced in a bout of offline brainstorming in which the culprits admit that they “were just literally trying to think of the most stupid company names”.

The site’s creators were inspired by what they consider to be the silly and contrived names of existing companies, such as *Accenture*, *Diageo* and *Consignia*. When interviewed by the BBC, the company described its brainstorming process as follows:

“We used an online Latin dictionary to come up with some of the names, and just added an 'i' or and 'a' to the end. Others, like *Ualeo*, we don't even know how to pronounce”

Names like *Ovisovis* suggest that a more general combinatorial process is also used. So this is a brainstorming process not unlike the one described here, in which reusable word fragments – catchy Latinate roots and their allowable (and cool-sounding) affixes – are mixed and matched to create novel combinations. The company does this for very different reasons, of course, since *WhatBrandAreYou.com* is intended to showcase the perils of simplistic and rule-governed design, as evident in the expensive but uninspired names crafted by rival consultants. Nonetheless, the BBC reports that some of the spoof names, such as *Tempero*, *Integriti*, *Winwin* and *Ovisovis*, have subsequently been snapped up and registered as legal company names by businesspeople with fewer design qualms. Of these ironic successes *The Design Conspiracy* observes: “clearly, we have an

aptitude for thinking up company names. But then, it's a lot easier than it seems”.

The web-site *WhatBrandAreYou.com* really is a design conspiracy then, one that shows that the combinatorial generation of new words is more akin to *brainstem*-storming than brainstorming. Yet, in spite of themselves, even these attempts at trivialization cannot help but produce words and names that other people find useful and creative. As mockumentary star David St. Hubbins comes to realize in *This is Spinal Tap*, “It's such a fine line, between stupid and clever”. Though brainstorming is always a scattershot process, we shall attempt to nudge the process toward the clever side of this fine line. Firstly, we will not generate purely random word mash-ups, by fusing together any prefix with any suffix; and secondly, we will not judge a candidate to be successful based on how stupid it sounds, but on how meaningful and useful it appears.

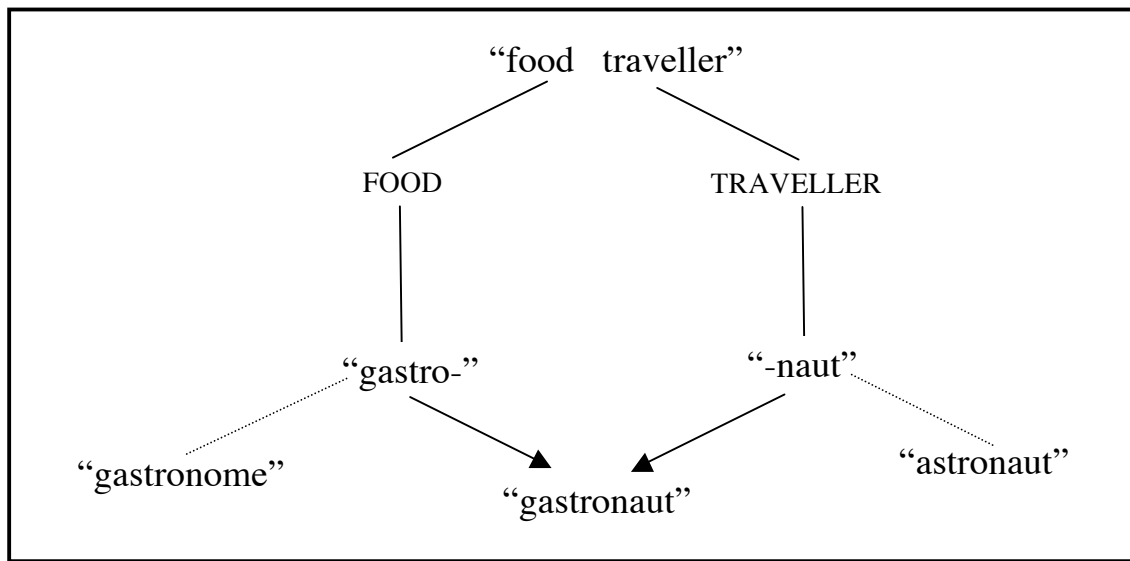


Figure 3. A common phrase (in a given context) provides the starting point for lexical invention. Word fragments like “gastro-“ have a conventional meaning in our inventory, but also suggest the words (and their meanings) in which they are recognizably used.

We do this by viewing word generation as a goal-oriented process. New words are coined, after all, in response to specific linguistic needs, so we begin by imagining what these needs might be. Suppose we find ourselves using the phrase “food traveller” so often, with connotations that arise from neither of the words “food” or “traveller” (such

as SELECTIVE and ADVENTUROUS), that we wish to coin a catchy new label for these footloose gourmands. With this goal in mind, we can now consider all of the possible ways of expressing the two component ideas, FOOD and TRAVELLER, as compatible word fragments that can be fused into an integrated whole. Looking to our inventory of fragments, we see that the prefix “gastro-“ can signify FOOD and that the suffix “-naut” can signify TRAVELLER. Furthermore, “-naut” is most recognizably used in the word “astronaut”, with has the connotations of bravery and adventure that we seek in our new creation. Putting both fragments together, we arrive at the word “*gastronaut*”; this simple process is illustrated in Figure 3.

Though a simple process, this is neither random nor haphazard. Note how the word fragments “gastro-“ and “-naut” are suggestive of the words they are most associated with, so the new coinage “*gastronaut*” can also be seen as a portmanteau of sorts, of the words “*gastronome*” and “*astronaut*”. Now imagine replicating this simple process on a much grander scale. Using a large text collection in which we can find brainstorming fodder like “*food traveller*”, we identify those compound descriptions that can give rise to a potentially useful fusion like “*gastronaut*”, by mapping the individual words into their corresponding prefix and suffix fragments. As our brainstorming fodder, we can use the *Google n-grams*, a huge collection of short English text snippets from Google’s search index with a frequency of 40 or more hits on the web. Running through every two-word snippet in Google’s 2-grams with our inventory of 400-plus word fragments, we generate many more than 150 new words; in fact, we generate over 90,000 meaningful candidates.

Because we use meaningful phrases as our starting point, these candidates are never truly meaningless, no matter how odd they may sound. So the best criteria for ranking these candidates do not touch on their apparent stupidity, but on the following issues:

1. Does the word appear in a conventional dictionary (i.e., is it an existing word)?
2. Does the word already have a web presence (e.g., with a web frequency ≥ 40)?

Criterion (1) considers whether brainstorming has thrown up a word that everyone can agree is *not* a novel coinage; to maintain our focus on new words, we ignore here the real possibility that brainstorming has assigned a creative new meaning to an existing word.

Criterion (2) considers whether the word, if not already established, is nonetheless already in use on the web; this covers words that reside in the grey area between popular acceptance by web users and official acceptance by print dictionaries. If either of these criteria holds, then the word cannot be considered especially creative, though it can be considered quite meaningful. Conversely, if neither of these criteria holds, then the word may not be an acceptable coinage (despite its potential to be meaningful) but it does have the potential to be considered truly creative. This is the essential trade-off of creative exploration, for it is only when we go beyond the marked territories on our maps and charts that we open ourselves to the possibilities (and risks) of real discovery.

Following the assessment criteria proposed by Margaret Boden, words that obey criterion (1) are not creative, while those that obey (2) but not (1) are original with respect to the person or machine that produces them. These words are thus likely to be *P-Creative* (or *Psychologically-Creative*), since they are novel and meaningful to their creator, if not always new to the larger community. Words that obey neither (1) nor (2) have the potential to be, in Boden's classification scheme, truly *H-Creative* (*Historically-Creative*), since no discernible record exists of these words having being created before. Of course, as *The Design Conspiracy* would eagerly point out, there may be good reasons why these potentially *H-Creative* words have never before been coined by a sensible person whose goal was not to deliberately "think of the most stupid names". Creativity may well happen in private, but the proof of *H-Creativity* (versus *H-Stupidity*) ultimately lies in the recognition and adoption of these private efforts by the wider society.

The PH-Test

As St. Hubbins observed, the line between clever and stupid is remarkably fine, and most new words thrown up by the brainstorming process cluster ambiguously around this line; are they *H-Creative*, or merely *H-Stupid*? Of the 90,409 words that are suggested, just 769 words (a mere 0.85%) can be found in a conventional dictionary (we used WordNet, an online electronic dictionary and thesaurus, for this test; bigger dictionaries are unlikely to raise this percentage significantly above 1%). Another 2,690 words (or 3%) are found to be in relatively common usage on the web (we used the Google 1-grams database for

this test, and words must have a web frequency of 40 or more to be listed as a 1-gram). The remainder, constituting a full 96% of the words that are generated, are unlikely to be found in any print dictionary, and have no significant presence of the web. Like many of the designer objects in modern society, they may be very clever or very stupid, and in some cases they can even be both.

We can sample this collection of potentially creative words by looking at all of the words that are brainstormed around the same word fragment. Let's focus here on the suffix "naut", which earlier gave us the interesting *P-Creative* coinage "gastronaut". Brainstorming gives us a wider range of uses for a fragment like "naut" than one might expect, and the size of the following list speaks for itself. In this list, dictionary-defined words are shown in bold, while *P-Creative* uses with a web-frequency of 40 or higher are underlined; words with at least one usage on the web (at the time of writing) are shown in italics. Note that just one meaning is listed for each term, although some terms can be suggested by several meanings, and different terms can arise from the same meaning.

Gerontonaut	"age explorer"	<i>Aironaut</i>	"air traveller"
Oxionaut	"air traveller"	Aeronaut	"air traveller"
<i>Avionaut</i>	"airplane traveller"	<i>Xenonaut</i>	"alien explorer"
Allelonaut	"alternative traveller"	<i>Paleonaut</i>	"ancient explorer"
<i>Archeonaut</i>	"history explorer"	<i>Anthonaut</i>	"archive explorer"
<i>Taxonaut</i>	"type explorer"	<u><i>Spacernaut</i></u>	"space explorer"
<i>Phononaut</i>	"audio explorer"	<i>Autonaut</i>	"automatic explorer"
<i>Aristonaut</i>	"royal traveller"	<i>Optinaut</i>	"great traveller"
<i>Primonaut</i>	"first explorer"	<i>Magninaut</i>	"big explorer"
<u><i>Bionaut</i></u>	"life explorer"	<i>Biblionaut</i>	"library explorer"
<u><i>Neuronaut</i></u>	"mind explorer"	<u><i>Psychonaut</i></u>	"crazy traveller"
<i>Pontonaut</i>	"bridge explorer"	<i>Stegonaut</i>	"hidden traveller"
Numismatonaut	"cash traveller"	<i>Aetionaut</i>	"cause explorer"
<i>Speleonaut</i>	"cave explorer"	<i>Cytenaut</i>	"cell explorer"
<i>Dynonaut</i>	"change explorer"	<i>Arterionaut</i>	"channel explorer"
<i>Vianaut</i>	"portal explorer"	<i>Rheonaut</i>	"circulation explorer"
<i>Civinaut</i>	"citizen explorer"	<i>Citinaut</i>	"city explorer"
<i>Metronaut</i>	"city traveller"	<i>Juxtanaut</i>	"close traveller"
<i>Genenaut</i>	"gene explorer"	<i>Cryptonaut</i>	"code explorer"
<i>Spectronaut</i>	"light explorer"	<i>Chloronaut</i>	"green traveller"

<i>Chromatonaut</i>	“color explorer”	<i>Holonaut</i>	“complete traveller”
<i>Encyclonaut</i>	“comprehensive traveller”	<i>Meronaut</i>	“component explorer”
<i>Zeugmanaut</i>	“connection explorer”	<i>Geonaut</i>	“world traveller”
<i>Gastronaut</i>	“gourmet traveller”	<i>Idionaut</i>	“personal traveller”
<i>Nyctonaut</i>	“night traveller”	<i>Plutonaut</i>	“money traveller”
<i>Oceanonaut</i>	“ocean explorer”	<i>Heteronaut</i>	“diverse traveller”
<i>Diplonaut</i>	“diplomatic traveller”	<i>Disconaut</i>	“disk explorer”
<i>Telenaut</i>	“distant traveller”	<i>Cynonaut</i>	“dog traveller”
<i>Pictonaut</i>	“image explorer”	<i>Oneironaut</i>	“dream explorer”
<i>Aquilonaut</i>	“eagle explorer”	<i>Protonaut</i>	“first traveller”
<i>Terranaut</i>	“earth explorer”	<i>Electronaut</i>	“electronic traveller”
<i>Ergonaut</i>	“energy explorer”	<i>Thermonaut</i>	“energy explorer”
<i>Photonaut</i>	“light explorer”	<i>Anglonaut</i>	“English explorer”
<i>Histerionaut</i>	“time explorer”	<i>Phraseonaut</i>	“expression explorer”
<i>Ophthalmonaut</i>	“eye explorer”	<i>Visionaut</i>	“picture explorer”
<i>Videonaut</i>	“image explorer”	<i>Tachonaut</i>	“fast traveller”
<i>Matronaut</i>	“female traveller”	<i>Oligonaut</i>	“powerful traveller”
<i>Pyronaut</i>	“fire explorer”	<i>Bromonaut</i>	“food explorer”
<i>Ludonaut</i>	“fun traveller”	<i>Nymphonaut</i>	“girl traveller”
<i>Bononaut</i>	“good traveller”	<i>Chironaut</i>	“hand explorer”
<i>Misonaut</i>	“hate explorer”	<i>Chrononaut</i>	“time traveller”
<i>Anthroponaut</i>	“man explorer”	<i>Cryonaut</i>	“ice explorer”
<i>Infonaut</i>	“message explorer”	<i>Radionaut</i>	“radio explorer”
<i>Endonaut</i>	“inner explorer”	<i>Hibernonaut</i>	“winter traveller”
<i>Veronaut</i>	“knowledge explorer”	<i>Limnonaut</i>	“lake explorer”
<i>Lexinaut</i>	“word explorer”	<i>Verbonaut</i>	“language traveller”
<i>Macronaut</i>	“large traveller”	<i>Patronaut</i>	“male traveller”
<i>Cartonaut</i>	“map explorer”	<i>Maginaut</i>	“master traveller”
<i>Glossonaut</i>	“message explorer”	<i>Mnemonaut</i>	“memory explorer”
<i>Logiconaut</i>	“science traveller”	<i>Psychonaut</i>	“soul explorer”
<i>Lunonaut</i>	“moon traveller”	<i>Cinenaut</i>	“movie explorer”
<i>Audionaut</i>	“sound explorer”	<i>Econaut</i>	“nature explorer”
<i>Neonaut</i>	“new traveller”	<i>Noxinaut</i>	“night traveller”
<i>Nomonaut</i>	“number explorer”	<i>Oleonaut</i>	“oil explorer”
<i>Mononaut</i>	“single traveller”	<i>Typonaut</i>	“print explorer”
<i>Phytonaut</i>	“plant explorer”	<i>Floranaut</i>	“plant explorer”
<i>Tectonaut</i>	“structure explorer”	<i>Dystonaut</i>	“poor traveller”
<i>Temponaut</i>	“time explorer”	<i>Spironaut</i>	“ring explorer”
<i>Robonaut</i>	“robotic explorer”	<i>Lithonaut</i>	“rock explorer”
<i>Petronaut</i>	“rock explorer”	<i>Quixonaut</i>	“romantic traveller”

<i>Spherenaut</i>	“world traveller”	<i>Technonaut</i>	“science explorer”
<i>Veneranaut</i>	“sex traveller”	<i>Eronaut</i>	“sex explorer”
<i>Conchonaut</i>	“shell explorer”	<i>Dermanaut</i>	“skin explorer”
<i>Cielonaut</i>	“sky traveller”	<i>Micronaut</i>	“tiny explorer”
<i>Leptonaut</i>	“tiny traveller”	Astronaut	“space explorer”
<i>Acronaut</i>	“tall traveller”	<i>Teleonaut</i>	“task explorer”
<i>Horonaut</i>	“time traveller”	Aquanaut	“water explorer”
<i>Nanonaut</i>	“tiny explorer”	<i>Chemonaut</i>	“chemical explorer”
<i>Hydronaut</i>	“water explorer”	Meteoronaut	“weather explorer”
<i>Oenonaut</i>	“wine explorer”	Dipsonaut	“wine traveller”
<i>Logonaut</i>	“word explorer”	Cosmonaut	“universe explorer”

In all, brainstorming produces 142 uses for the suffix “-naut” alone. Just four of these are found in a standard print dictionary, while another fifteen are found in relatively common usage on the web. Yet brainstorming from meanings to words ensures that most of the remaining candidates seem both meaningful and interesting, and many even seem useful, if perhaps only humorously so. In fact, 95 of these others are found at least once on the web (and are shown in italics, without underlining or bolding, above). A single usage on the web is a very low quality bar indeed, but such usages at least suggest that the terms are well-formed phonetically and that someone, somewhere, has used them to mean *something*. At the time of writing, for instance, only one web document contains the word “Oleonaut”, which was brainstormed to mean “an oil explorer”. Within this web document – a thread about biodiesel in an online forum – “*Oleonaut*” is the chosen pseudonym of a poster who is a self-described worker “in a major oil company refinery”. Similarly, the term “*Paleonaut*” has but a single web usage that can be found by Google, a forum posting that alludes to the “archaeological tasks” that one must be willing to perform to upgrade old software that is no longer supported. The brainstormed meaning “ancient explorer” is clearly apt here, if a little over-stated, humorously describing anyone who must explore the crumbling ruins of old and decrepit legacy software.

Licence To Drill

These terms and their glosses are not the final output of the brainstorming process, but an intermediate product that can lead to other meanings and insights. After all, glosses like

“oil explorer” and “ancient explorer” are not meanings *per se*, but more-or-less ambiguous phrases that simply suggest meanings, and figuring out these precise meanings and their consequences in a particular context is a key part of the brainstorming process. The individual words of these glosses need not be interpreted according to their dominant dictionary senses, but can be interpreted with their own measure of creativity. So the “ancient” of “ancient explorer” is not the “ancient” of “ancient Babylon” but the “ancient” of curmudgeonly old men, 8-track tapes and black-and-white TV sets. Likewise, the glosses for both “*Oleonaut*” and “*Paleonaut*” use “explorer” in a metaphorical rather than a literal sense, not as one who pushes back the boundaries of geographic knowledge, but as one who “explores options and possibilities”.

This potential for metaphorical interpretation means that simple glosses can often serve as a springboard for further creative brainstorming. Consider the gloss “*royal traveller*” for “*Aristonaut*”. The description “royal” can be taken literally here, to denote an aristocratic jet-setter who flies in style while circling the globe on high-profile diplomatic visits. Alternatively, it can be taken metaphorically, to describe the entitled attitude of ultra-rich passengers who only ever fly first-class. Indeed, it is easy to imagine the label “*Aristonaut*” actually being used by an airline to describe its own customers in the first-class cabin, or perhaps to describe those frequent flyers whose accumulated air-miles put them beyond the gold and even platinum level of entitlement, to the full Aristocratic level (passengers on *Air France* might actually enjoy the irony). Some web-users aptly employ the term to describe the billionaires who buy their way into space – so-called “space tourists” or self-funded astronauts – and who usually stay at that orbiting space hotel, the International Space Station (ISS). Since one would have to be filthy rich to buy a commercial ticket into orbit, the novel term “*Plutonaut*” also seems appropriate here, though neither “*Aristonaut*” nor “*Plutonaut*” have gained widespread acceptance in this descriptive role.

Interestingly, if one favours the “space tourist” interpretation of “*Aristonaut*” and “*Plutonaut*” then this makes each term a fully-fledged portmanteau, for “*Aristocratic astronaut*” and “*Plutocratic astronaut*” respectively. Conversely, if one prefers the “*royal traveller*” and “*cash traveller*” interpretations, to describe a well-heeled traveller (and not an astronaut *per se*), then the same terms no longer function as a portmanteau fusion of

existing words. A portmanteau is not just a word with a particular form, but a word that appears to have been formed by a particular process to suggest a particular meaning. Yet when the words that contribute to a portmanteau are themselves a composite of meaningful word-fragments, such as “aristo-“ or “-naut”, the end result can appear the same: a word that can be viewed both as a juxtaposition of reusable components (to suggest only a juxtaposition of those component ideas) and as a fusion of two clipped words (to suggest a full integration of the meanings of both whole words, and not just their clipped parts). Even simple brainstorming of the kind described here can yield words whose meaning *and* form is subject to further creative manipulation.

Flux Capacitor Not Included

H. G. Wells coined the term “Time Machine” in his 1895 book of the same name, and it has since become the preferred name for any device that allows travellers to move back and forth in time. Wells was not the first writer to explore the notion of time travel; in fact, he wasn’t even the first to think of time travel in mechanical terms, facilitated by a machine rather than by magic, dreams or some other mysterious plot device. A “time machine” is a somewhat uncreative name for a machine that allows time travel (it could just as well describe a clock), but it certainly makes for a catchy title, and what it lacks in creativity it makes up for in simplicity and resonance. From a Don Norman design perspective it has good visibility and good mapping, and tells the audience everything it needs to know about the story inside. Wells presumably reasoned that the concept of time travel was a creative enough premise as it was (if not exactly a *H-Creative* one), and already challenging enough for his audience, without additionally saddling his story with an obscure and challenging name like “*The Chronomat*”. But Wells was not always a fan of simple and unpretentious names. Seven years earlier, when publishing a precursor to his time travel adventure in a school journal, he gave it the enigmatic if not entirely successful title “*The Chronic Argonauts*”. Fortunately, good sense prevailed when the time came to write the novel, and he eventually ditched the label, if not the idea, of a “*chronic Argonaut*”. Today, though time travel is still confined to the realms of fiction and speculative physics, the conceit is more popular than ever, and as if to mirror the simplicity of “time machine” the phrase “time traveller” is still the most common label

for fictional characters who jump through time. Just as some people have faces that are made for radio, “chronic Argonaut” is a clunky name whose charms are definitely more conceptual than linguistic. Surely we can brainstorm a more creative label for this notion?

Our simple but rather broad-ranging brainstorming exercise throws up a variety of one-word labels that may fit the bill. For instance, a “*Horonaut*” (clock explorer) might aptly describe time travellers who make small jumps in time, as measured on a clock rather than on a calendar. Alternately, the terms “*Histerionaut*” and “*Archeonaut*” might well describe a time traveller who makes significant leaps back into history, while a “*Paleonaut*” might describe really adventurous travellers who jump all the back into prehistoric times. In physics, a tachyon is a particle that can travel faster than the speed of light and thereby move backwards in time, so a “*Tachonaut*” (fast traveller) might likewise describe a time traveller who exploits the same temporal loopholes. More generally, the terms “*Chrononaut*” and “*Temponaut*” are just compressed ways of saying “*time traveller*”, though “*Temponaut*” does suggest that time travellers have (or need) a good sense of rhythm. As we saw earlier, we can arrive at the same neologism via different paths. The word “Chrononaut” can be seen both as a simple combination of “chrono-” (time) with “-naut” (traveller) *and* as the portmanteau product of clipping “Chronic” and “Argonaut” so as to squeeze the parts into a single streamlined word. If Wells’ earlier title seems half-baked then perhaps it really is just half-formed, an ugly duckling that wants to grow up to be a graceful portmanteau swan.

More Half-Baked Words

Though it generates a large number of novel and interesting candidates, our brainstorming exercise is nonetheless limited by its reliance on a fixed inventory of word fragments with pre-determined meanings. This inventory has combinatorial power, but it means that the word-generation process is not free to combine an open-ended diversity of different ideas. Yet there is a simple solution to this dilemma, one that broadens the scope of the generation process while still ensuring that only meaningful word elements are purposefully combined.

Just as the fragment “astro-” can (loosely) mean SPACE when used as a word prefix,

so can the word “space” itself, as evidenced by words such as “*spacenaut*”. Obviously, any stand-alone word can suggest its own conventional meaning when it is used as a fragment of a larger word, so we can easily extend our inventory of word elements to include any whole word that can combine with one of the existing suffixes or prefixes already in our inventory. The meaning gloss associated with this new element will be the word itself, so the assigned meaning of “space” is SPACE, the meaning of “tourist” is TOURIST, and so on. To avoid over-stocking the inventory with words that we really should not be thrusting into combination with others, we can limit ourselves to words that have an attested use as part of a larger form, in combination with one of our standard inventory elements. For instance, the Google 1-gram “*spacenaut*” attests to the validity of “space” as a prefix that means SPACE because the combined meaning “*space explorer*” is also an attested phrase in the Google 2-grams. Just as some single women prefer married men because somebody somewhere has already found them relation-worthy, we will only look to words that have proven themselves worthy of *P-Creative* combination in the past. So, the 1-gram “*astrotourist*” attests to the validity of adding “tourist” as a suffix meaning TOURIST, since the combined meaning “*space tourist*” is also an attested 2-gram. When expanded with these word elements and their rather obvious meanings, our inventory is capable of generating a much larger space of candidate neologisms; though many of these combinations will now seem half-baked, in form if not in meaning, this familiarity may actually make them easier to digest by their target audience.

A *P-Creative* coinage like “astrotourist” (used on the web, but not yet listed in a dictionary like WordNet) thus paves the way for a potentially *H-Creative* coinage like “*pharmatourist*” or “*biotourist*”, since “tourist” now becomes a freely combinable suffix in our inventory of word parts. Similarly, “*ecoterrorist*” marks out “-terrorist” as a viable suffix, leading to “*infoterrorist*”, “*technoterrorist*” and “*gastroterrorist*” (none of which is an attested Google 1-gram). If the late TV chef Keith Floyd was the first self-described “gastronaut”, perhaps Gordon Ramsay and Marco Pierre White are suitable donkeys on which to pin the tail marked “*gastroterrorist*”? Clearly the set of words that can be used to supply parts for larger words (what linguistics call “free morphemes”, or meaning elements that are free to either stand alone or bond with others) is considerably larger than the set of word fragments that can be used only to accessorize other words (linguists

call these “bound morphemes”, to reflect their status as the indentured bondsmen of the lexical world). The latter is a more-or-less *closed* set, one that we made a decent attempt to inventory as part of our earlier brainstorming exercise, while the former is an *open* set, since almost any content word can be used as part of a longer and more complex form.

Unsurprisingly, our trawl through the Google 1-grams finds thousands of words that have attested uses as parts of larger forms. Words like “man”, “web”, “tech” and “blog” are commonly used as suffixes in various meaning-stuffed confections (such as “*astroman*”, “*biblioweb*”, “*neurotech*” and “*photoblog*”), while words like “zoo”, “sex” and “echo” are just as promiscuous as prefixes (in words such as “*zoopedia*”, “*sexorama*” and “*echotrope*”). In all, the 1-grams yields a haul of over different 7000 content words that have been used with a prefix from our inventory of parts, and almost 2000 words that have been used with a suffix from this inventory. When these new elements are added to our inventory and the brainstorming exercise is re-run – now with the added possibility of producing half-baked as well as fully-cooked neologisms – we see a massive increase in the number of novel candidates that are generated. Once again, only words that are motivated by pre-existing phrases (with hopefully sensible meanings) are considered, such as “*metroblog*” for “*urban blog*” and “*narcotourist*” for “*drug tourist*”. Surprisingly, perhaps, brainstorming now generates over a million such words. While few will be to everyone’s taste and many will forever be the lexically unborn, this is a remarkably large space of potential words, one that dwarfs the listings of even the largest dictionaries.

The new words are the lexical equivalent of test-tube babies: created *in vitro* in the lab, but not always viable *in vivo*. Both processes can be wasteful because both are subject to risks, so many more possibilities are created than will ever survive or thrive. The combinatorial process at the heart of brainstorming ramps up word production to an almost industrial scale, so that potentials that are evidenced in just a small number of words *in vivo* are magnified and catalyzed into a much larger set of possibilities *in vitro*. Consider the word “snob”: it has a nasty little meaning but a delightfully compact sound. We find just one Google 1-gram in which “snob” is used as a suffix in combination with one of our inventoried prefixes: “*vinosnob*”, meaning a “*wine snob*”. But this one instance is enough to suggest that “snob” has a much greater potential as a word suffix,

and when “-snob” is added to our inventory as a suffix (meaning SNOB, naturally), brainstorming cranks out 95 other candidates from this one meagre example, including:

Girosnob	“art snob”	Phonosnob	“sound snob”
Magnisnob	“big snob”	Brontosnob	“huge snob”
Maxisnob	“biggest snob”	Bibliosnob	“book snob”
Logosnob	“word snob”	Mobilesnob	“car snob”
Chocosnob	“chocolate snob”	Cinesnob	“film snob”
Metrosnob	“city snob”	Taxosnob	“class snob”
Cryptosnob	“code snob”	Encyclosnob	“complete snob”
Holosnob	“total snob”	Gastrosnob	“culinary snob”
Oleosnob	“fat snob”	Ethnosnob	“culture snob”
Choreosnob	“dance snob”	Cynosnob	“dog snob”
Ichtyosnob	“fish snob”	Bromasnob	“food snob”
Ludosnob	“game snob”	Chronosnob	“history snob”
Archeosnob	“old snob”	Endosnob	“inner snob”
Lexisnob	“language snob”	Dipsosnob	“liquor snob”
Audiosnob	“music snob”	Paleosnob	“old snob”
Monosnob	“only snob”	Phytosnob	“plant snob”
Lithosnob	“rock snob”	Aristosnob	“royal snob”
Cryosnob	“snow snob”	Civisnob	“society snob”
Techosnob	“technology snob”	Bellisnob	“war snob”
Aquasnob	“water snob”	Cyclosnob	“wheel snob”

There are as many kinds of snobs as there are things to be elitist about in modern society. If some of the above possibilities seem odd or implausible, like “*snow snob*”, remember that the meaning in each case (the two-word gloss in quotes) is an attested phrase from the Google 2-grams. A moment’s thought suggests that a “*snow snob*” is a skiing enthusiast with annoyingly superior ideas about what constitutes the best kind of snow – or the best slopes – for skiing. In some cases the non-availability of a given phrase in the Google 2-grams means that a potential word meaning is not found by the brainstorming process. For instance, the phrase “*secret snob*” is not an attested 2-gram, so brainstorming overlooks the possibility that a “*Cryptosnob*” might be a “*secret snob*”, or

in other words, a person who keeps their “*Endosnob*” well hidden. In contrast, “*secret tourist*” is an attested 2-gram (with a web-frequency of 75), which gives us the word “*cryptotourist*” for travellers who sneak into territories without the necessary permissions.

It’s Not What You Know, But the Way That You “New” It

As even this simple exercise in brainstorming demonstrates, *in vitro* experiments that use the right stock of ingredients can cause our test-tubes to runneth over with many plausible candidates for meaningful new words. Whether these test-tube creations take root *in vivo* is a matter of context, fashion, and personal need. Our creative urge to coin new words is often driven by the need to assign a convenient and meaningful handle to a composite idea that we wish to manipulate as though it were a single concept. By assigning a single atomic symbol to this complex of ideas, we encourage and reinforce the practice of viewing these ideas as a single coherent chunk of knowledge. In these cases, the creative naming process is motivated by the realization that a given piece of knowledge is important enough, or useful enough, to have its own name. But the naming process can also work in the opposite direction: the fact that a given piece of knowledge can be reduced to a single well-formed atomic word suggests that this knowledge is coherent enough to be useful. So rather than assign a name *after* an idea has proven itself interesting, we can use the viability of a name as a *prior* indicator of which ideas in a sea of competing ideas are most worthy of our limited attention span. After all, the space of new and meaningful words is governed by linguistic conventions and intuitions, and can be navigated using past linguistic experience – as captured in a large text corpus – as a guide. In essence, we can use the corpus-guided brainstorming of interesting new words as a convenient and well-structured proxy for the brainstorming of interesting new ideas.

A rose by any other name might smell just as sweet, but roses wouldn’t seem half as poetic if they were saddled with a stinker of a name like “bloodwort” or “thornweed”. Indeed, if we failed to linguistically discriminate roses from other kinds of flower, they would lack not just a distinct name but any poetic use or resonance. Shakespeare’s warning is a linguistically astute one, and we should always be careful to not confuse a word with its meaning or its referent, but names *are* important, and designer products

deserve designer names. Names give structure to our thoughts, and facilitate the crystallization of our vague impressions into a firm concept, or at least a concept that feels firm enough to grasp and manipulate. Or as Goethe put it, “When ideas fail, words come in very handy”. Sometimes a name conveys more solidity than is conceptually present, giving rise to empty jargon that obfuscates and disguises a speaker’s ignorance to the detriment of an idea and its audience. But even in these cases a meaningful name can act as a conceptual skeleton on which more flesh can later be hung by more sincere speakers. Names help us construct our shared ontology of reality, to organize what we know, and by creative implication, the innovative possibilities that we don’t yet know. In the much-quoted words of Donald Rumsfeld, erstwhile sage of the Pentagon:

“There are known knowns. These are things we know that we know. There are known unknowns. That is to say, there are things that we know we don’t know. But there are also unknown unknowns. There are things we don’t know we don’t know.”

Rumsfeld’s ruminations are, surprisingly, as relevant to the workings of linguistic creativity as they are to military intelligence. For one, Rumsfeld actually coins a new word here: we are all familiar with the noun use of “unknown” and its plural “unknowns” to denote the missing elements of a logic puzzle or mathematical equation that one must solve for (in the latter, unknowns are named by variables such as “X”); but Rumsfeld uses figure-ground reversal to name the negative space around “unknowns” as “*knowns*”, a form whose novelty causes spell-checkers everywhere to break out the red ink. This is a subtle form of creativity that works so well in this context – as a natural companion to the familiar “unknown” – that it hardly draws attention to itself as a lexical innovation. For another, Rumsfeld’s ontology explicitly distinguishes the different kinds of knowledge that are implicated in linguistic creativity. All creativity leverages existing knowledge to reach further insights; these “*known knowns*” are vital to creativity, but they are not themselves creative. Goal-directed creativity seeks to find answers to pressing questions that demand a solution; these questions concern “*known unknowns*” and can be addressed using convergent means (narrowly focussed deduction) or Guilford-style divergent means (wide-ranging speculation), with the latter more likely to throw up a creative solution.

Linguistic creativity is often playfully gratuitous, used more for showboating than for clarity. In these cases the creative urge becomes the creative splurge, as authors mischievously use creativity to disguise “*known knowns*” to make them appear more like “*known unknowns*”. In their efforts to avoid cliché, for instance, authors sometimes coin an obviously novel form (a “*known unknown*”) that has precisely the same meaning as a popular cliché (one of many culturally entrenched “*known knowns*”), safe in the knowledge that the audience can decipher their intent by retrieving the appropriate cliché. In other words, these creative authors get to keep their cliché and beat it too. In the following 2009 article from the Irish Independent daily newspaper, the columnist Kevin Myers coins the dubious term “*The Dubble*” as a portmanteau of both “*Dublin Bubble*” and “*Dublin Hubble*”, to represent the Irish media’s views on (and from) the Irish economic bubble:

“The Dubble has a range of about three miles, namely the radius of the Dublin Bubble, from the Montrose Tower to Liberty Hall”

Having established this lexical analogy with the Hubble telescope, Myers goes on to say:

“Dubble was at its umbiliscopic capers again on Sunday [November 8th, 2009]”

The coinage “*umbiliscopic*” seems puzzling at first, and will strike most readers as a lexical “*known unknown*”. Clearly it has some relevance to the telescope analogy established earlier, but what exact meaning does it convey? A moment’s reflection reveals that the term is a fancy workaround for the cliché “*navel gazing*”, since the prefix “umbili-” is derived from “umbilicus” which in turn is a clinical term for the “navel” or “belly button” (though most of us will arrive at this meaning via the word “umbilical”), and the suffix “-scopic” refers to any instrument for viewing or observing and thus for “gazing”. This is a combination that can easily be generated by our earlier brainstorming exercise, using as it does word fragments with established glosses that together yield a common phrase/meaning (“umbili-” + “-scopic” = “navel gazing”). Myers gets to develop his telescope conceit and sneak in a cliché too, but in a gratuitous yet creative form that momentarily challenges his audience. At the time of writing Google indexes just two web documents with the word “umbiliscopic”: one is Myers’ article, the other a

medical text that refers to a treatment called “*umbiliscopic cholecystectomy* (UC)”, the surgical removal of the gallbladder via an incision in the navel.

Rumsfeld’s ontology is incomplete since he declines to recognize the implied but seemingly oxymoronic fourth possibility of “*unknown knowns*”. This is the tacit knowledge that allows us to see acts of linguistic creation as novel enough to seem innovative (“*unknowns*”) but familiar enough to be understood via existing meanings (“*knowns*”). Rumsfeld’s own coinage of “*knowns*” relies on “*unknown knowns*”: the word form is new and initially unknown, but its meaning is, in a sense, already known, because it exhibits the visible and natural mapping (in Don Norman’s design philosophy) between form and meaning that allows us to turn an adjective into a noun. Likewise, Myers’ “*umbiliscopic*” is a novel (i.e., previously “*unknown*”) means of conveying a conventional (i.e., widely “*known*”) phrase or meaning. Creativity relies on “*known knowns*” but often takes place amongst the “*known unknowns*” and, more daringly, amongst the “*unknown unknowns*”. Naming is a crucial part of this creative trajectory from unknown to known.

In 2003, Retired Admiral John Poindexter advocated the development of a futures market in *terror* as part of the Pentagon’s ill-fated *Total Information Awareness* (TIA) initiative. It was scrapped by Rumsfeld, as was Poindexter, when the scheme became a “*known known*” in the media. Intriguingly, this market would have allowed the Pentagon to track potential terrorist activity by encouraging investors to profit financially from correct predictions about future terrorist attacks, assassinations and regime overthrows. The events of 9/11 showed that terrorists are at their most dangerous when they are creative, and what better way to anticipate future creativity than by soliciting as broad a divergence of views as possible, and letting the profit motive decide which candidates are most likely? A market forces us to attach meaningful symbolic names to individual possibilities, rather like company names in a stock exchange, and so gives these possibilities a minimal conceptual substance that can be grasped and further elaborated by users of the market. In fact, a marketplace provides the divergence of possibilities that J. P. Guilford argued is central to the generation of creative alternatives, while also providing a dynamic and informed basis for converging on the most attractive candidates.

Suppose we wanted to brainstorm the potential forms that a terrorist outrage might take. It would certainly help if we could survey the divergent types of terrorist and terrorism that we (and our text corpora) are familiar with, to construct a taxonomy of possible scenarios. Looking to our simple brainstorming exercise with words, which earlier gave us 96 types of “snob”, we find that “terrorist” is even more productive as a suffix, yielding over 200 lexical variations. While forms like “*aeroterrorist*” (“*air terrorist*” or “*flying terrorist*”), “*chemoterrorist*” (“*chemical terrorist*”) and “*xenoterrorist*” (“*foreign terrorist*” or “*alien terrorist*”) suggest that this exercise is akin to shutting the barn door after the pale horse has bolted, it is instructive to be reminded of the chilling diversity that terrorism can assume. Besides, while the names that we brainstorm are plausible and grounded in the familiar, their meanings are creatively under-specified and open to further elaboration. The phrase “*ancient terrorist*” prompts the neologism “*archeoterrorist*”, which can denote terrorists in ancient times, aging terrorists in the present time, or archeological terrorists who target ancient sites of cultural value (such as the Coliseum in Rome, the Pyramids of Giza, and so on). Other forms suggest that animal terrorists (“*faunaterrorist*”), sound terrorists (“*audioterrorist*”), food terrorists (“*gastroterrorist*”), drug terrorists (“*narcoterrorist*”), book terrorists (“*biblioterrorist*”) and mind terrorists (“*psychoterrorist*”) are also atypical possibilities whose potential meanings are worthy of further creative investigation.

But it is not all doom and gloom: the suffix “tourist” is more productive still, yielding almost 250 new forms. One can imagine a tour company exploring new marketing opportunities by brainstorming different types of tourist for the 21st century, from dream tourists and mind tourists (“*oneirotourist*” and “*psychotourist*” respectively) who never leave their living rooms, to farm tourists (“*agrotourist*”), fast tourists (“*tachotourist*”), fire tourists (“*pyrotourist*”) and game tourists (“*ludotourist*”). In each case, it is the possibility of coining a creative and meaningful new word that allows us to eke out the hidden potentials of the often unlikely phrases in our huge text corpora.

Conclusions: The Two Donalds

In this chapter we’ve looked at the process of coining creative new words from two

unusual but complementary perspectives. The first perspective uses Donald A. Norman's design philosophy to view new words as consumer goods that must compete in the linguistic marketplace. For new words to not just survive, but actually thrive in common parlance, they must exhibit all the best features of a well-designed piece of hardware: that is, their surface forms must map naturally and visibly onto their conceptual meanings. The second perspective comes from a second and altogether more controversial Donald, one famed for his love – if not execution – of radical ideas. In the ontological distinctions drawn by Donald Rumsfeld, acquired knowledge is “known” while as-yet-un-acquired knowledge is “unknown”. New words often strike us as “*known unknowns*” – forms we know we don't know – yet these words are typically not random or arbitrary: a thoughtful creator has combined “known” elements to create a word that can, with basic insight, be revealed as an unknown (novel) variation on a known (familiar) theme.

We began this chapter by considering the creative principles at work in phrasal names like *Virgin Mary* and *Bloody Shame*, before then focusing our energies on designer words, lexical creations that squeeze multiple design elements into a single word form. For instance, these sham drinks and their non-alcoholic brethren now commonly go by the label “mocktail”, a portmanteau concoction of “mock” and “cocktail” that could not do a better job of pushing customers toward the comforts of a real drink. The best designer words are simultaneously new *and* familiar, combining known elements into novel and superficially unknown forms.