OPINION

Diagnosing the decline in pharmaceutical R&D efficiency

Jack W. Scannell, Alex Blanckley, Helen Boldon and Brian Warrington

Abstract | The past 60 years have seen huge advances in many of the scientific, technological and managerial factors that should tend to raise the efficiency of commercial drug research and development (R&D). Yet the number of new drugs approved per billion US dollars spent on R&D has halved roughly every 9 years since 1950, falling around 80-fold in inflation-adjusted terms. There have been many proposed solutions to the problem of declining R&D efficiency. However, their apparent lack of impact so far and the contrast between improving inputs and declining output in terms of the number of new drugs make it sensible to ask whether the underlying problems have been correctly diagnosed. Here, we discuss four factors that we consider to be primary causes, which we call the 'better than the Beatles' problem; the 'cautious regulator' problem; the 'throw money at it' tendency; and the 'basic research–brute force' bias. Our aim is to provoke a more systematic analysis of the causes of the decline in R&D efficiency.

Over the past 60 years, there have been major advances in many of the scientific and technological inputs into drug research and development (R&D). For example, combinatorial chemistry increased the number of drug-like molecules that could be synthesized per chemist per year by perhaps 800-fold during the 1980s and 1990s[1–3], and greatly increased the size of chemical libraries[4]. DNA sequencing has become over a billion times faster since the first genome sequence was determined in the 1970s[5–7], aiding the identification of new drug targets. It now takes at least three orders of magnitude fewer man-hours to calculate three-dimensional protein structure via X-ray crystallography than it did 50 years ago[8,9], and databases of three-dimensional protein structure have 300 times more entries than they did 25 years ago[9] (see the RCSB Protein Data Bank database website), facilitating the identification of improved lead compounds through structure-guided strategies. High-throughput screening (HTS) has resulted in a tenfold reduction in the cost of testing compound libraries against protein targets since the mid-1990s[10]. Added to this are new inventions (such as the entire field of biotechnology, computational drug design and screening, and transgenic mice) and advances in scientific knowledge (such as an understanding of disease mechanisms, new drug targets, biomarkers and surrogate end points).

There have also been substantial efforts to understand and improve the management of the commercial R&D process. Experience has accumulated on why projects overrun[11], on the factors that influence financial returns on R&D investment[12–17], on project success[18] and R&D portfolio management[19–22], on how to reduce costs by outsourcing, and on what is likely to impress or worry the regulatory authorities[23].

However, in parallel — as many have discussed — R&D efficiency, measured simply in terms of the number of new drugs brought to market by the global biotechnology and pharmaceutical industries per billion US dollars of R&D spending, has declined fairly steadily[24]. We call this trend 'Eroom's Law', in contrast to the more familiar Moore's Law ('Eroom's Law' is 'Moore's Law' backwards). Moore's Law is a term that was coined to describe the exponential increase in the number of transistors that can be placed at a reasonable cost onto an integrated circuit. This number doubled every 2 years from the 1970s to 2010. The term is used more generally for technologies that improve exponentially over time. The data in FIG. 1a show that the number of new US Food and Drug Administration (FDA)-approved drugs per billion US dollars of R&D spending in the drug industry has halved approximately every 9 years since 1950, in inflation-adjusted terms. Part of the contrast between Moore's Law and Eroom's Law is related to the complexity and limited current understanding of biological systems versus the relative simplicity and higher level of understanding of solid-state physics[25] but, as discussed below, there are other important causes.

Although there are difficulties in making like-for-like comparisons in R&D spending over very long periods, Eroom's Law has been fairly robust. The number of new drugs introduced per year has been broadly flat over the period since the 1950s, and costs have grown fairly steadily[24]. The slope of the line, over 10-year periods at least, does not change substantially (FIG. 1b), and assumptions about the delay between R&D investment and drug approval do not have a substantial influence on the overall pattern (FIG. 1c). For more details of the data used for FIG. 1, and the major assumptions made, see Supplementary information S1 (table).

Eroom's Law indicates that powerful forces have outweighed scientific, technical and managerial improvements over the past 60 years, and/or that some of the improvements have been less 'improving' than commonly thought. The more positive anyone is about the past several decades of progress, the more negative they should be about the strength of countervailing forces. If someone is optimistic about the prospects for R&D today, they presumably believe the countervailing forces — whatever they are — are starting to abate, or that there has been a sudden and unprecedented acceleration in scientific, technological or managerial progress that will soon become visible in new drug approvals.

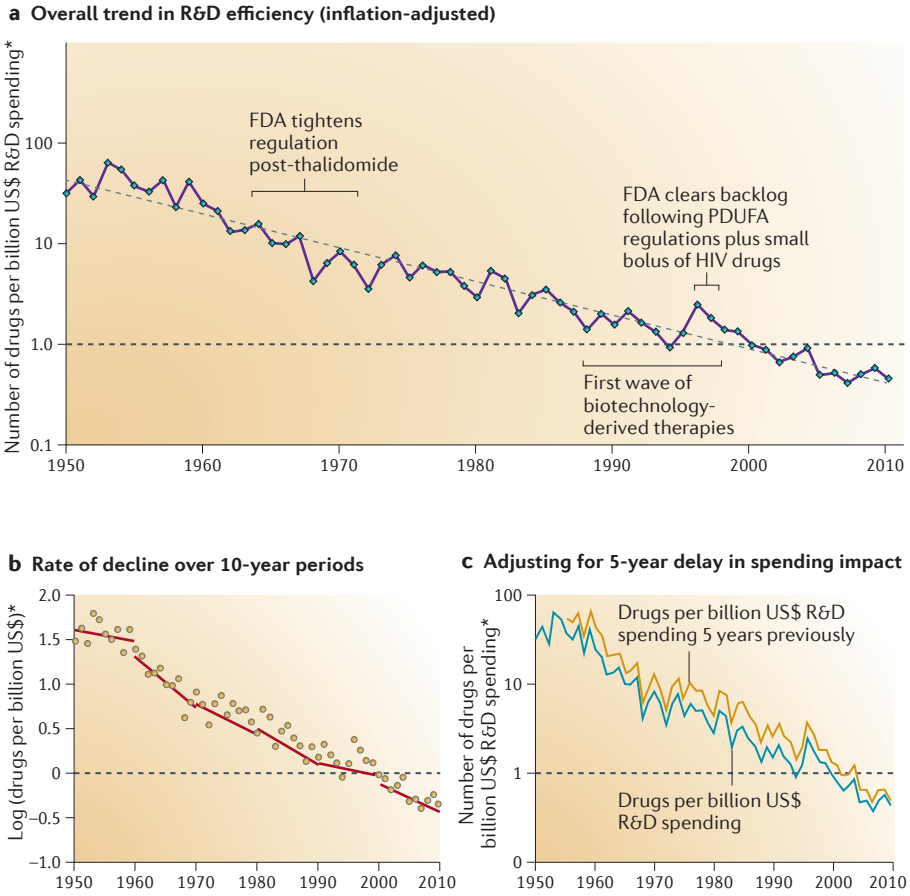

Figure 1 | **Eroom's Law in pharmaceutical R&D. a** | The number of new drugs approved by the US Food and Drug Administration (FDA) per billion US dollars (inflation-adjusted) spent on research and development (R&D) has halved roughly every 9 years. **b** | The rate of decline in the approval of new drugs per billion US dollars spent is fairly similar over different 10-year periods. **c** | The pattern is robust to different assumptions about average delay between R&D spending and drug approval. For details of the data and the main assumptions, see Supplementary information S1 (table) and REFS 24,86,87. Note that R&D costs are based on the Pharmaceutical Research and Manufacturers of America (PhRMA) Annual Survey 2011 (REF. 86) and REF. 87. PhRMA is a trade association that does not include all drug and biotechnology companies, so the PhRMA figure understates R&D spending at an industry level. The total industry expenditure since 2004 has been 30–40% higher than the PhRMA members' total expenditure, which formed the basis of this figure. The new drug count, however, is the total number of new molecular entities and new biologics (applying the same definition as Munos[24]) approved by the US FDA from all sources, not just PhRMA members. We have estimated real-term R&D cost inflation figures from REFS 24,87. The overall picture seems to be fairly robust to the precise details of cost and inflation calculations. Panel **a** is based on a figure that originally appeared in a Bernstein Research report (The Long View — R&D productivity; 30 Sep 2010). *Adjusted for inflation. PDUFA, Prescription Drug User Fee Act.

The magnitude and duration of Eroom's Law also suggests that a lot of the things that have been proposed to address the R&D productivity problem are likely, at best, to have a weak effect. Suppose that we found that it cost 80 times more in real terms to extract a tonne of coal from the ground today than it did 60 years ago, despite improvements in mining machinery and in the ability of geologists to find coal deposits. We might expect coal industry experts and executives to provide explanations along the following lines: "The opencast deposits have been exhausted and the industry is left with thin seams that are a long way below the ground in areas that are prone to flooding and collapse." Given this analysis, people could probably agree that continued investment would be justified by the realistic prospect of either massive improvements in mining technology or large rises in fuel prices. If neither was likely, it would make financial sense to do less digging.

However, readers of much of what has been written about R&D productivity in the drug industry might be left with the impression that Eroom's Law can simply be reversed by strategies such as greater management attention to factors such as project costs and speed of implementation[26], by reorganizing R&D structures into smaller focused units in some cases[27] or larger units with superior economies of scale in others[28], by outsourcing to lower-cost countries[26], by adjusting management metrics and introducing R&D 'performance scorecards'[29], or by somehow making scientists more 'entrepreneurial'[30,31]. In our view, these changes might help at the margins but it feels as though most are not addressing the core of the productivity problem.

There have been serious attempts to describe the countervailing forces or to understand which improvements have been real and which have been illusory. However, such publications have been relatively rare. They include: the FDA's 'Critical Path Initiative'[23]; a series of prescient papers by Horrobin[32–34], arguing that bottom-up science has been a disappointing distraction; an article by Ruffolo[35] focused mainly on regulatory and organizational barriers; a history of the rise and fall of medical innovation in the twentieth century by Le Fanu[36]; an analysis of the organizational challenges in biotechnology innovation by Pisano[37]; critiques by Young[38] and by Hopkins *et al.*[39], of the view that high-affinity binding of a single target by a lead compound is the best place from which to start the R&D process; an analysis by Pammolli *et al.*[19], looking at changes in the mix of projects in 'easy' versus 'difficult' therapeutic areas; some broad-ranging work by Munos[24]; as well as a handful of other publications.

There is also a problem of scope. If we compare the analyses from the FDA[23], Garnier[27], Horrobin[32–34], Ruffolo[35], Le Fanu[36], Pisano[37], Young[38] and Pammolli *et al.*[19], there is limited overlap. In many cases, the different sources blame none of the same countervailing forces. This suggests that a more integrated explanation is required.

Seeking such an explanation is important because Eroom's Law — if it holds — has very unpleasant consequences. Indeed, financial markets already appear to believe in Eroom's Law, or something similar to it, and the impact is being seen in cost-cutting measures implemented by major drug companies. Drug stock prices indicate that investors expect the financial returns on current and future R&D investments to be below the cost of capital at an industry level[40], and

would prefer less R&D and higher dividends. Investors may well be wrong about this. However, they have less reason to be biased towards optimism about the likelihood of Eroom's Law being successfully counteracted than those who are working in the industry, or those who sell consulting services to the industry. Shareholders ultimately appoint executives and control resource allocation, so their perceptions matter. It is likely that Pfizer, Merck & Co., AstraZeneca and Eli Lilly will be spending less — in nominal terms — in 2015 than they did in 2011, partly in response to shareholder pressure. Across the top ten large pharmaceutical companies, it seems that nominal R&D spending will be flat until 2015, which represents a decline in real terms. More importantly, the combined effect of declining real-term R&D spending with Eroom's Law (fewer new drugs per billion US dollars of R&D investment over time) is that there will be fewer new drugs and/or drugs will become inordinately expensive. This will threaten the huge benefits[41,42] that follow from the availability of effective and affordable new drugs.

In our view, avoiding such an outcome requires a more systematic analysis of the factors that underlie Eroom's Law. We think that any serious attempt to explain Eroom's Law should try to address at least two things: the progressive nature of the decline in the number of new drugs per billion US dollars of R&D spending, and the scale (~80-fold) of the decline. In this article, we make some suggestions. We realize that the industry is heterogeneous, so our generalizations will be wrong in many cases. We appreciate the intellectual effort that has been made by others on analysing the problems of R&D productivity. We note that our chosen measure of R&D efficiency is based on cost per new drug approved. This does not account for the huge variation in the medical and financial value of new drugs. A few breakthrough drugs — for example, a highly effective Alzheimer's disease treatment — could have much greater medical and financial value than a larger number of new drugs that provide only modest incremental benefits. We also note that the very long cycle time for drug R&D means that our productivity measure is a lagging indicator; perhaps things have improved, but the result is not yet visible.

However, with the aim of prompting debate and analysis, here we discuss what we consider to be the four primary causes of Eroom's Law: the 'better than the Beatles' problem; the 'cautious regulator' problem; the 'throw money at it' tendency; and the 'basic research–brute force' bias. There may also be some contribution from a fifth factor, termed 'the low-hanging fruit' problem, but we consider this to be less important.

Primary causes

The 'better than the Beatles' problem. Imagine how hard it would be to achieve commercial success with new pop songs if any new song had to be better than the Beatles, if the entire Beatles catalogue was available for free, and if people did not get bored with old Beatles records. We suggest something similar applies to the discovery and development of new drugs. Yesterday's blockbuster is today's generic. An ever-improving back catalogue of approved medicines increases the complexity of the development process for new drugs, and raises the evidential hurdles for approval, adoption and reimbursement. It deters R&D in some areas, crowds R&D activity into hard-to-treat diseases and reduces the economic value of as-yet-undiscovered drugs. The problem is progressive and intractable.

Few other industries suffer from this problem. In the example of the coal industry noted above, the opencast deposits are mined first. However, the coal is burnt, which increases the value of the coal that is still in the ground. In most intellectual property businesses (for example, fashion or publishing), people get bored with last year's creations, which maintains demand for novelty. Unfortunately for the drug industry, doctors are not likely to start prescribing branded diabetes drugs because they are bored with generic metformin.

Anti-ulcerants — still a very valuable therapeutic area in terms of revenues — provide an example of the shadow that is cast by successful drugs. A class of anti-acid agents known as potassium-competitive acid blockers, such as soraprazan (now discontinued), would probably be safe and effective anti-ulcerants, and 15 years ago they could have been blockbusters. The problem today is that there are now two classes of highly effective and safe anti-ulcer drugs on the market: the histamine H_2 receptor antagonists (for example, generic ranitidine, which is available over the counter) and the proton pump inhibitors (for example, generic esomeprazole and several others). Any sensible healthcare system would only pay for patients to receive a new branded potassium-competitive acid blocker if they fail to respond to a cheap generic proton pump inhibitor and/or H_2 receptor antagonist, but such patients are a very small proportion of the total population. This general problem applies in diabetes, hypertension, cholesterol management and many other indications.

Pammolli *et al.*[19] have provided a quantitative illustration of the 'better than the Beatles' problem. Their analysis compared R&D projects started between 1990 and 1999 with those started between 2000 and 2004. Attrition rates rose during the latter period. However, the increase could be largely explained by a shift in the mix of R&D projects from commercially crowded therapeutic areas in which historic drug approval probabilities were high (for example, genitourinary drugs and sex hormones) to less crowded areas with lower historical approval probabilities (for example, antineoplastics and immunomodulatory agents).

There is another related potential cause of Eroom's Law that has frequently been put forward, termed the 'low-hanging fruit' problem, which results from the progressive exploitation of drug targets that are more technically tractable[43]. To be clear, the 'low-hanging fruit' problem argues that the easy-to-pick fruit has gone, whereas the 'better than the Beatles' problem argues that the fruit that has been picked reduces the value of the fruit that is left in the tree.

In our opinion, the 'low-hanging fruit' problem is less important than the 'better than the Beatles' problem. First, estimates of the number of potential drug targets[44,45] versus the number of drugged targets[46] suggest that many decades-worth of new targets remain if the industry continues to exploit four or five new targets each year. It is also becoming clear that many drugs may derive their therapeutic benefit from interactions with multiple proteins rather than a single target. These drugs are 'magic shotguns' rather than 'magic bullets'[47]. If this turns out to be more generally true, then worrying about the 'low-hanging fruit' problem would be similar to worrying that a shortage of notes is threatening the future of music composition. In our view, the 'low-hanging fruit' explanation is sometimes tautological as 'technically easy' tends to be equated with 'already discovered'[48]. Indeed, investigation of the history of drug discovery suggests that there was little easy or obvious about some of the great discoveries of the 1940s and 1950s, such as the anti-inflammatory effects of corticosteroids, the psychiatric effects of imipramine or lithium, or the antibacterial properties of penicillin[36,49–51].

The 'cautious regulator' problem. Progressive lowering of the risk tolerance of drug regulatory agencies obviously raises the bar for

the introduction of new drugs, and could substantially increase the associated costs of R&D[52]. Each real or perceived sin by the industry, or genuine drug misfortune, leads to a tightening of the regulatory ratchet, and the ratchet is rarely loosened, even if it seems as though this could be achieved without causing significant risk to drug safety. For example, the Ames test for mutagenicity may be a vestigial regulatory requirement; it probably adds little to drug safety but kills some drug candidates. Furthermore, for most of the past 60 years large and sophisticated drug companies may not have been disappointed to see the regulatory ratchet tighten because it reduced competition.

It also seems that the concern that drug companies could cheat the system in some way has led the cautious regulator to apply an audit-based approach to regulatory documentation, as the more demanding the reporting requirements are, the harder it is to cheat without leaving some kind of error or inconsistency in what is reported. The scale of reporting was summarized recently by the Chief Scientific Officer of Novo Nordisk in the company's third quarter 2011 results conference call with respect to the submission to the FDA of data on two new insulin therapies: "If printed and stacked, the many million pages of documentation, with a total of 9 million electronic links, [would] exceed the height of [the] Empire State Building."

The impact of the 'cautious regulator' problem on Eroom's Law is apparent in FIG. 1. First, it shows R&D efficiency dipping in the early 1960s following the 1962 Kefauver Harris Amendment to the Federal Food, Drug, and Cosmetic Act, which was introduced in the wake of the thalidomide drug safety disaster. For the first time, medicines had to demonstrate efficacy, and the safety hurdles were also raised. This reduced financial returns on R&D for a decade or so[12,14], before rising drug prices outstripped R&D cost inflation and increased financial returns in the 1970s[15]. Interestingly, FIG. 1 also shows a rise in R&D efficiency in the mid to late 1990s, which is likely to be due to two regulatory factors: primarily the clearing of a regulatory backlog at the FDA following the implementation of the 1992 Prescription Drug User Fee Act (PDUFA), but also a small contribution from the rapid development and approval of several HIV drugs. In the case of HIV drugs, organized and politically astute lobbying effectively lowered the normal regulatory hurdles[53].

The 'cautious regulator' problem follows, in part, from the 'better than the Beatles' problem, as the regulator is more risk-tolerant when few good treatment options exist; or, put another way, the availability of safe and effective drugs to treat a given disease raises the regulatory bar for other drugs for the same indication. Although the 'cautious regulator' problem is tractable in principle, it is hard to see the regulatory environment relaxing to any great extent. Society may be right to prefer a tougher regulator, even if it means more costly R&D. Drug safety matters. And although the 1950s and 1960s may be viewed by some as a golden age in terms of therapeutic innovation[36,48,54], it seems unlikely that the severe adverse outcomes for many patients taking part in clinical trials during this period[36] would be acceptable today.

The 'throw money at it' tendency. The 'throw money at it' tendency is the tendency to add human resources and other resources to R&D, which — until recent years — has generally led to a rise in R&D spending in major companies, and for the industry overall. It is probably due to several factors, including good returns on investment in R&D for most of the past 60 years, as well as a poorly understood and stochastic innovation process that has long pay-off periods. In addition, intense competition between marketed drugs (where being second or third to launch is often worth less than being first) provides a rationale for investing additional resources to be the first to launch. There may also be a bias in large companies to equate professional success with the size of one's budget.

Unfortunately for people working in R&D today, tackling the 'throw money at it' tendency looks feasible. Investors and many senior executives think that a lot of costs can be cut from R&D without reducing output substantially. Whether this is correct remains to be seen, although if so, it may be the single strategy most likely to counteract Eroom's Law in the short term. The risk, however, is that the lack of understanding of factors affecting return on R&D investment that contributed to relatively indiscriminate spending during the good times could mean that cost cutting is similarly indiscriminate. Costs may go down, without resulting in a substantial increase in productivity.

The 'basic research–brute force' bias. The 'basic research–brute force' bias is the tendency to overestimate the ability of advances in basic research (particularly in molecular biology) and brute force screening methods (embodied in the first few steps of the standard discovery and preclinical research process) to increase the probability that a molecule will be safe and effective in clinical trials (FIG. 2). We suspect that this has been the intellectual basis for a move away from older and perhaps more productive methods for identifying drug candidates[32–34]. It should be noted that many of our comments are more relevant to small-molecule drugs, although the data used for FIG. 1 also include biologics.

FIGURE 2 illustrates the standard model of most drug R&D. It is — effectively — a serial search, filter and selection process. Scientific and technical advances have, superficially at least, increased the breadth of the funnel (that is, more potential targets have been identified, and more drug-like molecules have been synthesized). They have improved the filtering efficiency by several orders of magnitude (for example, HTS versus testing in expensive and low-throughput animal models). They should also have increased the quality of filtering and selection (for example, the use of pathway analysis for target selection, the use of transgenic mice for target validation and the use of accumulated experience to favour molecules that would be likely to have good ADMET (absorption, distribution, metabolism, excretion and toxicology) characteristics).

The cumulative effect of improvements in these early steps should have been a higher signal-to-noise ratio among drug candidates entering clinical trials; that is, the candidates chosen should have had a greater likelihood of successfully demonstrating effectiveness and safety in these trials. This, in turn, should have increased R&D efficiency, given that most of the costs of new drug development are related to the costs of failed projects[22]. Yet the probability that a small-molecule drug successfully completes clinical trials has remained more or less constant for 50 years[21], and overall R&D efficiency has declined[24].

So how can some parts of a process improve dramatically, yet important measures of overall performance remain flat or decline? There are several possible explanations, but it seems reasonable to wonder whether companies industrialized the wrong set of activities[34,36,38]. At first sight, R&D was more efficient several decades ago (FIG. 1), when many research activities that are today regarded as critical (for example, the derivation of genomics-based drug targets and HTS) had not been invented, and when other activities (for example, clinical science, animal-based screens and iterative medicinal chemistry) dominated.

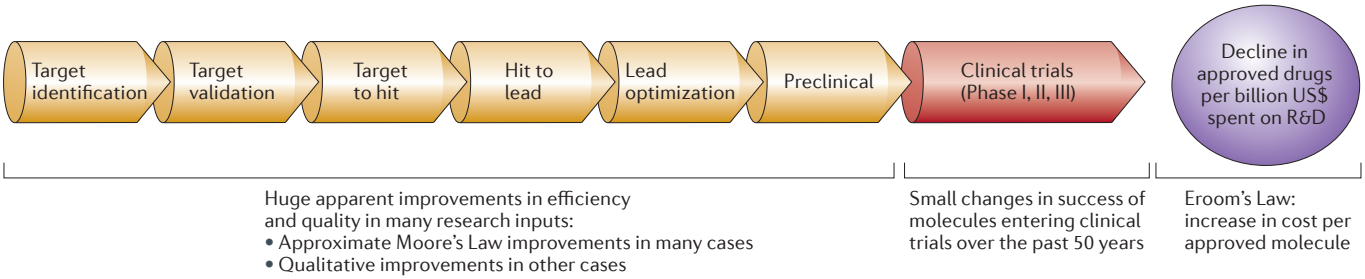

Figure 2 | **How can some parts of the R&D process improve, yet the overall efficiency decline?** Dramatic improvements in brute force screening methods and basic science should have tended to increase the efficiency of the research process (more leads tested against more targets, at a lower cost; shown in gold) and raised its quality (better targets as disease pathways and mechanisms are understood, better leads that avoid old mistakes surrounding ADMET (absorption, distribution, metabolism, excretion and toxicity) characteristics, and so on). This, in turn, should have increased the probability that molecules would succeed in the clinic (shown in red), which in turn should have increased overall efficiency, as research and development (R&D) costs are dominated by the cost of failure. However, the probability that a small molecule successfully completes clinical trials has remained more or less constant for 50 years[21], whereas overall R&D efficiency has declined[24]. One possible explanation for this is that much of the industry industrialized and 'optimized' the wrong set of R&D activities.

There have been several interesting critiques of modern research[33,48,55], but here we highlight two potential problems. First, much of the pharmaceutical industry's R&D is now based on the idea that high-affinity binding to a single biological target linked to a disease will lead to medical benefit in humans[39]. However, if the causal link between single targets and disease states is weaker than commonly thought[38,56], or if drugs rarely act on a single target, one can understand why the molecules that have been delivered by this research strategy into clinical development may not necessarily be more likely to succeed than those in earlier periods.

Indeed, drug-like small molecules tend to bind promiscuously, and this sometimes turns out to have an important role in their efficacy[47,57] as well as their so-called off-target effects[39]. Targets are parts of complex networks leading to unpredictable effects[58], and biological systems show a high degree of redundancy, which could blunt the effects of highly targeted drugs[56,57]. Perhaps this helps to explain why the R&D process was more cost-effective several decades ago (FIG. 2), when expensive labour-intensive animal models — rather than cheap automated molecular assays — formed the basis of initial drug screening[36,49–51,59].

More recent analysis also points to a similar conclusion. More first-in-class small-molecule drugs approved between 1999 and 2008 were discovered using phenotypic assays than using target-based assays[60]. Drugs approved during this period would have been discovered when screening activity was dominated by the target-based approach, so one might have expected more target-based discoveries. Perhaps target-based approaches are efficient for pursuing validated therapeutic hypotheses but are less effective in the search for innovative drugs that have a better chance of clearing the 'better than the Beatles' barrier.

The second potential problem follows from the nature of chemical space and a shift from iterative medicinal chemistry coupled with parallel assays (pre-1990s) to serial filtering that begins with HTS of a static compound library against a target. Directed iteration — even if each cycle is slow — may be a much more efficient way of searching a large and high-dimensional chemical space than fast HTS of a predefined collection of compounds (BOX 1).

As an aside, biologics have had a higher success rate than small molecules once they leave research and enter clinical trials. There was an approximately 32% approval rate for biologics versus an approximately 13% approval rate for small-molecule drugs first tested in humans between 1993 and 2004 (REF. 21). This may not be surprising for copies or close analogues of endogenous signalling molecules (for example, insulins, erythropoietins or growth hormones) or for agents that replace dysfunctional proteins (for example, clotting factors, lysosomal enzymes, and so on). The high rates of success in clinical trials of monoclonal antibodies (and related fusion proteins) is perhaps more notable[61]. One might expect them to suffer from the same kind of problems with single-target efficacy as small molecules (albeit with fewer off-target effects). However, they have opened up new sets of therapeutic targets, which may suffer less from the 'better than the Beatles' problem. Perhaps their success is also a function of their limited target set — either cell surface proteins or protein-based extracellular signalling molecules. In both cases, the chain of causality between target binding and therapeutic effect is relatively short. Out of 34 monoclonal antibodies or other targeted biologics (such as fusion proteins or aptamers) that have been approved by the FDA, 13 target white blood cell-specific antigens (for example, CD20) and are used for haematological cancers or immunosuppression; three target receptors in the human epidermal growth factor receptor family and are used in oncology; seven target tumour necrosis factor or interleukins and are used for immunomodulation in autoimmune diseases; and four target vascular endothelial growth factor variants and are used in oncology or ophthalmology.

In our view, there are several reasons why the 'basic research–brute force' bias has come to dominate drug research. First, by the early 1980s there was already a sense that the pace of pharmaceutical innovation was slowing. The 'cautious regulator' problem was an obvious drag[52,54,62]. The 'better than the Beatles' problem was starting to emerge, with complaints that new drugs had only modest incremental benefit over what was already available[62]. There were concerns about the 'low-hanging fruit' problem, with a growing sense that the industry had started to run out of good animal models to screen drugs for still poorly treated diseases[52,62].

Second, the 'basic research–brute force' bias matched the scientific zeitgeist[48], particularly as the older approaches for early-stage drug R&D seemed to be yielding less. What might be called 'molecular reductionism' has become the dominant stream in biology in general, and not just in the drug

Box 1 | **Directions in small-molecule drug discovery**

The 1990s saw a major shift in small-molecule drug discovery strategies, from iterative low-throughput *in vivo* screening and medicinal chemistry optimization to target-based high-throughput screening (HTS) of large compound libraries. At first sight, the former is slow and expensive in terms of the number of compounds that can be tested, whereas the latter is fast and cheap[59]. However, the topography of chemical space and the nature of industrialized drug discovery may conspire to make the second approach less productive. The problem is not necessarily HTS per se (the pros and cons of which are actively debated[79]); rather, it may be the research processes that new technologies helped to cement.

First, real-world compound libraries for HTS cover infinitesimally small and somewhat redundant regions of chemical space, which is vast; it has been suggested that there could be between 10^{26} and 10^{62} (REFS 80,81) chemotypes that would comply with the Lipinski guidelines for oral drugs[82], and each chemotype has a large number of potential derivatives. By contrast, a typical corporate screening collection for HTS contains around 10^6 chemical entities and perhaps 10^3 chemotypes. Furthermore, mergers have revealed that different companies' compound libraries often substantially overlap. This is not surprising: companies generated their libraries in similar ways, as they used clustered sets of molecules from similar historical campaigns; there is a limited set of commercially available reagents; and a relatively small number of reactions are amenable to high-throughput automated synthesis.

Second, it has proved to be difficult to design systems that reward people for producing 'good' hits and leads rather than 'more' hits and leads. Collections are biased towards developable compounds with acceptable ADME (absorption, distribution, metabolism and excretion) characteristics. Companies want measurable developability benchmarks. There are few immediate prizes for chemical or biological novelty. The pre-selection and pre-design of screening collections means that the lead structures are largely foreseen. It provides no easy way to jump from local chemical optima to something better.

Third, the process to whittle down a few thousand HTS hits into a couple of qualified leads has been dominated by molecules that win on potency measures. Selection is based on serial assays, with most molecules failing at each step. There is no practical way to view the full biological profile of all hits at an early stage. Hits with merely adequate target potency but with other potentially attractive features (such as good ADME, other interesting biological properties, and so on) could be thrown away. This further focuses the search process on small parts of screening collections. It may even focus the search process on a suboptimal part of the screening collection. Recent research suggests that there is a negative correlation between *in vitro* potency and desirable ADME and toxicology[83]. Given these features of HTS in the real world, we should expect different drug companies to produce similar molecules for a given target. We should also expect these molecules to reflect local optima within the screening collections, rather than global optima from the much larger chemical universe.

Before the 1990s, however, the standard approach for small-molecule drug discovery involved synthesizing and screening a relatively small number of compounds. There would be a few tens of molecules (often fewer) in active assessment at any one time, and perhaps 1,000 molecules synthesized by a team of chemists during a 5-year project. The search usually started with a molecule that was known, or suspected, to have promising pharmacology but perhaps with poor ADME characteristics: adrenaline led to the development of beta blockers, and histamine lead to the development of cimetidine. Phenomenological screening was also used, to a small extent, to provide starting points. Each molecule was then assessed in a range of concurrent assays (often *in vivo*[59], considering potency, ADME, toxicity, selectivity and so on). Molecules were then modified (or discarded) depending on the results of the assays. The cycle was repeated, with the biological results being used to establish structure–activity relationships for each assay and thus advance the structures of lead compounds through the chemical space until one or two compounds met the multiple criteria necessary for progression into clinical trials. Unlike the screening case, after a few iterations one had compounds specifically customized to a particular target, with structures that would not have been foreseen at the start of the process. This approach prevented trial compounds from being confined to minor local optima. It facilitated what Sir James Black called "obliquity"[84] — the art of looking for one thing and finding something else. It made it less likely that competitors had identical drugs. Remarkably, the search for blockbuster drugs using this method was often achieved with fewer than 1,000 compounds.

This is a profoundly different search strategy to the one that was industrialized, but one that may be more efficient when there is a very large number of items arranged in a high-dimensional space, as is the case with drug-like molecules (see Supplementary information S2 (box)). This is because it is possible to traverse large regions of a high-dimensional space with a small number of steps[85], whereas any static, predefined compound library will cover only a tiny part of the chemical space. Perhaps this is part of the explanation of the pre-1990s productivity? These kinds of arguments are not lost on the drug industry. Efforts are underway to try to combine some of the obvious advantages of HTS with the advantages of small teams dedicated to a broader exploration of the biological profiles of a set of evolving lead compounds. The idea is to analyse several structure–activity relationships in parallel (for example, potency at the target, potency at likely toxicity sites, potency in cellular assays, *in vivo* ADME) to direct rapid, sometimes automated, iterative chemistry.

industry[33,34,55]: "Since the 1970s, nearly all avenues of biomedical research have led to the gene"[63]. Genetics and molecular biology are seen as providing the 'best' and most fundamental ways of understanding biological systems, and subsequently intervening in them[64]. The intellectual challenges of reductionism and its necessary synthesis (the '-omics') appear to be more attractive to many biomedical scientists than the messy empiricism of the older approaches.

Third, the 'basic research–brute force' bias matched the inclination of many commercial managers, management consultants and investors. The old model, based on iterative medicinal chemistry, animal-based screening and clinical science was seen as "too dependent on either inefficient trench-warfare type of slog or the unpredictable emergence of seemingly capricious geniuses like James Black, Paul Janssen, Daniel Bovet, Gertrude Elion, or Gerald Hitchings"[33]. Automation, systematization and process measurement have worked in other industries. Why let a team of chemists and biologists go on a trial and error-based search of indeterminable duration, when one could quickly and efficiently screen millions of leads against a genomics-derived target, and then simply repeat the same industrial process for the next target, and the next? In the early 1990s, few companies thought they could thrive or survive without moving towards a drug discovery process based on HTS and the products of the human genome.

Here, we are reminded of a debate[25] about improving clinical trial efficiency, triggered by an editorial by Andy Grove[65], the former Chief Executive of Intel — a man with personal experience of Moore's Law. Grove noted the "disappointing output" of R&D in the drug industry and made suggestions to radically change clinical trials by making more use of electronic health data[65]. Some biomedical scientists probably find Grove's intervention irritating, given the simplicity and predictability of semiconductor physics versus "biology's mysteries"[25]. However, shareholders and taxpayers have been persuaded to fund a lot of R&D because biomedical scientists (and drug industry executives) have told them that — thanks to molecular reductionism — it would soon become more predictable[63], more productive and less mysterious.

We think that the 'basic research–brute force' bias is supported by survivor bias among R&D projects. This makes drug discovery and development sound more prospectively rational than it really is. Nearly all

drugs are sold with a biological story that sounds like molecular reductionism and that sometimes, but not always, turns out to be true: for example, "drug x works by binding receptor a, which influences pathway b, which adjusts physiological process c, which alleviates disease d." Such stories get confused with prediction because we hear very little about the vast majority of the other projects that were also initiated on the basis of high-affinity binding of a plausible candidate to a plausible target, and that had similarly plausible biological stories until the point at which they failed in development for unexpected reasons.

It would be interesting to see how well prospective estimates of plausibility correlated with subsequent attrition. This point is illustrated by the anticancer drug iniparib. Attendees of the 2010 meeting of the American Society of Clinical Oncology (ASCO), or readers of the *New England Journal of Medicine*[66], could have been forgiven for believing that iniparib had a spectacular effect on metastatic breast cancer in a Phase II trial because it inhibited a specific target, poly(ADP-ribose) polymerase 1 (which is involved in DNA repair), and therefore potentiated chemotherapy. However, the following year, Phase III trial results presented at the 2011 ASCO meeting indicated that iniparib did not work very well in breast cancer[67], and it did not seem to inhibit poly(ADP-ribose) polymerase 1 very much either[68].

Fortunately, the 'basic research–brute force' issue is tractable in several ways. First, in a handful of therapeutic areas the research process does appear to be delivering better systems-level insights, better targets (or sets of targets) and better candidate drugs. Oncology is the most obvious example. It is hard to look at the genesis of drugs like crizotinib[69], vemurafenib[70] or vismodegib[71] and think that one is simply looking at random survivors. Furthermore, in oncology the regulator is less cautious and the back catalogue of approved drugs is far from 'Beatle-esque'. One or two other disease areas with simple genetics may perhaps resemble oncology. Second, more emphasis could be put on iterative approaches, on animal-based screening or even on early proof of clinical efficacy in humans, and less on the predictive power of high-affinity binding to the target of a molecule from a static library. Novartis is one company that is emphasizing proof-of-concept trials for drugs in rare diseases for which there is a high unmet need and a compelling match between the drug's mode of action and the disease. Only if there is success here does the company invest in more expensive trials in more common diseases in which the mode of action may be more speculative, or in which the risk–benefit profile may be less clear. Third, in some therapeutic areas people could just stop believing in the current predictive ability of 'basic research–brute force' screening approaches, and resist the temptation to put molecules into clinical trials without having more compelling evidence of the validity of the underlying therapeutic hypothesis.

There is, of course, no way of going back in time to see how well more recent R&D approaches would have worked in the 1940s and 1950s. It is possible that research has become much better at delivering the right molecules into the clinic but that the improvements have been swamped by the 'better than the Beatles' problem, the 'low-hanging fruit' problem and the 'cautious regulator' problem.

Ironically however, if the industry really has been doing the right things, the ultimate prognosis may be bleaker. One can think of the opportunities for R&D in terms of a Venn diagram: as science and technology improve, some sets grow (for example, the set of druggable targets, the set of drug-like molecules and the set of drugged targets), whereas other sets shrink (for example, the set of economically exploitable and still untreated diseases, or the set of acceptable off-target effects). It is obvious that R&D productivity could decline despite improvements in the inputs if the intersection that contained commercially attractive and approvable drug candidates shrunk. This idea is illustrated in FIG. 3, in which the notional set of validated targets grows between 1970 and 2010, but it does not grow fast enough to offset the growth in the set of targets that would either worry a cautious regulator or fail the 'better than the Beatles' test.

Finally, we note that it would be easier to improve the signal-to-noise ratio of drugs that enter clinical trials if: first, there was a detailed understanding of why drugs fail in the clinic; second, this led to the discovery of a small number of common failure modes; and third, this knowledge could be used to change the early stages of the R&D process. If it is impractical to carry out retrospective analyses on the precise molecular mechanisms of clinical trial failure, or if such retrospective analyses show that trials fail for many rare and idiosyncratic reasons, or if cycle times are so long that the lessons are obsolete by the time they are learned, then incremental improvement will be more difficult. Both the regulators[23] and the industry[18] are interested in the analysis of failure but it receives less scrutiny than one might expect given its dominant role in the costs of R&D.

Secondary symptoms

The four proposed primary causes of Eroom's Law discussed above have given rise to several 'symptoms' that tend to further increase costs, particularly the costs of clinical development. Some of these symptoms are highlighted below.

The narrow clinical search problem. The narrow clinical search problem is the shift from an approach that looked broadly for therapeutic potential in biologically active agents to one that seeks precise effects from molecules designed with a single drug target in mind. In the 1950s and 1960s, initial screening was typically performed in animals, not *in vitro* or *in silico*, and drug candidates were given in early stages of the development process to a range of physicians. Discovery involved, to an extent, the ability of physicians to spot patterns through careful clinical observation, especially in therapeutic areas in which symptomatic improvements are readily observable, such as psychiatry[36,49–51]. This is sometimes dismissed as serendipity but the approach made it likely that new therapeutic effects would be detected. Even recently, it appears that many — perhaps most — new therapeutic uses of drugs have been discovered by motivated and observant clinicians working with patients in the real world[72]. Some drug companies, particularly smaller and mid-sized firms, recognize this opportunity and are active repositioners of existing drugs.

However, the 'cautious regulator' problem and the 'basic research–brute force' bias have pushed most of the drug industry towards a narrow clinical search strategy. If a drug has an effect but this is not the precise effect that the trial designers anticipated, then the trial fails. Opportunities for serendipity are actively engineered out of the system. Perhaps it is too risky to let bright doctors with large numbers of patients make broad clinical observations, or to let creative scientists rummage around in rich clinical data sets, in case they find something unexpected, which has to be explained to the cautious regulator who then kills the project. Modern multicentre trials tend to spread the patients so thinly that a doctor who did want to look for patterns might miss them. In Phase II trials — perhaps the best opportunity to spot new things — the average number of patients

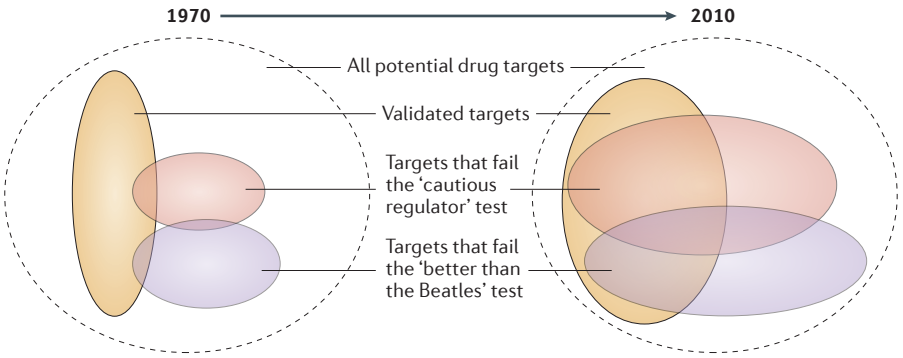

Figure 3 | **Venn diagram illustrating hypothetical headwinds to R&D efficiency.** Research and development (R&D) efficiency could decline if scientific, technical and managerial improvements are offset by other factors. For example, R&D efficiency could be limited by the supply of validated targets that could be drugged without failing the 'cautious regulator' test and/or the 'better than the Beatles' test. In this hypothetical illustration, the increase in the number of validated targets between 1970 and 2010 is outweighed by increasing regulatory caution and an improving catalogue of approved drugs.

per multicentre trial site is now very small: between five and ten patients in oncology, central nervous system and respiratory disease trials[73].

The big clinical trial problem. The first randomized controlled trial, published in 1948, recruited 109 patients and randomized 107 of them[74]. Between 1987 and 2001, the number of patients per pivotal trial for antihypertensive agents rose from around 200 to around 450 (REF. 75). Between 1993 and 2006, the average number of patients across the pivotal trials for a new oral antidiabetic drug rose from around 900 to over 4,000 (REF. 76). The first pivotal trial for Merck's simvastatin (a cholesterol-lowering agent), published in 1994, recruited around 4,400 patients[77]. A pivotal trial for Merck's anacetrapib, an investigational cholesterol-modulating agent intended to be used on top of drugs like simvastatin, is currently recruiting around 30,000 patients.

This expansion is a consequence of several factors. First, the 'better than the Beatles' problem increases trial size. Everything else being equal, clinical trial size should be inversely proportional to the square of the effect size. If the effect size halves, the trial has to recruit four times as many patients to have the same statistical power. The problem is that treatment effects on top of an already effective treatment are usually smaller than treatment effects versus placebo. Furthermore, Phase III trials have become a messy mixture of science, regulation, public relations and marketing. Trying to satisfy these multiple constraints tends to inflate their size and cost.

The best clinical trial to show efficacy would be something relatively small in a homogeneous patient sample recruited from as few centres as possible — the medical equivalent of a well-controlled experiment. But this tends to make the cautious regulator uneasy given variation in practice patterns and patients. What about rare side effects (the FDA has recently required post-marketing trials for long-acting bronchodilators in around 53,000 patients)? Small trials also make for bad marketing and, in the world of evidence-based medicine, poor market access. It is better to involve the senior doctors at the major centres. The number of principal investigators per drug in clinical trials has doubled over the past decade[73]. The consequence of this is multicentre trials that add noise and heterogeneity, and are therefore bigger and more expensive.

The multiple clinical trial problem. The 'better than the Beatles' problem has increased the complexity of medical practice. In some areas, where once there were only one or two treatment options, there is now a rich back catalogue. For example, the treatment of patients with type 2 diabetes was once a choice of insulin or diet and exercise, but can now involve a combination of drugs from around ten different drug classes: biguanides, thiazolidinediones, sulfonylureas, meglitinides, alpha-glucosidase inhibitors, dipeptidyl peptidase 4 inhibitors, glucagon-like peptide 1 analogues, amylin analogues, long-acting and short-acting insulin analogues, as well as various human insulins and insulin mixes. Treatment for patients with colon cancer was once a choice between surgical resection or palliative care, but now the National Comprehensive Cancer Network's colon cancer treatment guidelines contain up to 100 pages of detailed treatment algorithms.

The cautious regulator is less prepared to assume that the safety and efficacy of new drugs can be generalized across such heterogeneous and fragmented patient populations. Cost-sensitive health-care funders are also wary. This means narrower indications and more clinical trials per drug. The first long-acting insulin analogue, glargine, was approved by the FDA in 1999 following three pivotal Phase III trials. The newest long-acting insulin analogue, degludec, was filed for regulatory approval in 2011 following 12 pivotal trials (and, as mentioned above, an Empire State Building's worth of documentation). Some successful drugs in complex therapeutic areas appear to demand, over their life cycle, dozens of Phase III trials[78].

The long cycle time problem. In the 1950s and 1960s, cycle times were remarkably short by modern standards. The regulator was less cautious and there was less molecular reductionism before agents were screened for efficacy in animal models and in patients. This sped up innovation. The first antidepressant, imipramine, was synthesized in around 1951. It was screened almost immediately in rats, and tested personally by a few scientists at the drug company Geigy[51]. It was then tested without much success in various patient groups in 1952, tested again in 1953, found to be problematic in patients with psychosis in 1954 and tried yet again in 1955 before it was identified as an antidepressant in 1956. It completed preclinical development and had not one but three clinical cycles within 5 or 6 years. In 2005–2006, the typical period of time in clinical development for a new drug was over 9 years[21]. The biggest increase in development times came between the 1960s and the 1980s[21].

An idea: the CDDO

This article is intended to provoke further analysis of the forces that have countervailed scientific, technical and managerial improvements over the past 60 years. We have avoided cures, partly because the ratio of published cures to diagnoses is already too high. We do, however, have one idea, which might also be viewed as a thought experiment.

We suggest that all large drug companies introduce a new board level role, which we call the Chief Dead Drug Officer (CDDO). This role would be focused on drug failure

at all stages of R&D, and the CDDO would have a fixed time — for example, 18 months — from appointment to compose a detailed report that aims to explain the causes of Eroom's Law. This report would be submitted to the board of the company, included in the company's annual report to shareholders, and would also be submitted for publication in a scientific journal and sent to organizations such as the FDA and the US National Institutes of Health. The remuneration for the role would be structured in such a way as to provide a strong incentive to provide an accurate forecast of the future R&D productivity of the company and the industry overall. For example, perhaps the salary could be relatively modest, but the CDDO could be eligible for an enormous bonus if their projections after a 10-year period are no more than 10% too optimistic or no more than 30% too pessimistic.

We like the idea for several reasons. First, the CDDO has no incentive to be irrationally optimistic. Second, R&D costs are dominated by the cost of failure[73]. Most molecules fail. Most research scientists spend most of their time on products that fail. It seems fitting that someone on the board should focus on the products that consume most of the R&D organization's time, energy and money. Third, an expertise in drug failure should qualify the CDDO to produce a good explanation of Eroom's Law.

The CDDO's report should aim to explain the scale of the change in productivity. It should set out the major factors responsible for the progressive decline, and rank them in order of importance. It should consider how the relative importance of these factors has changed over time. Perhaps changes at the FDA dominated from 1960 to 1970, but something else dominates now? The analysis should compare different therapeutic areas. It should assess the extent to which the different factors are tractable. There should be some effort to quantify the 'better than the Beatles' problem and the 'low-hanging fruit' problem, as well as the potential value of underexploited drug targets. Attention should be given to the regulatory ratchet. Which requirements are most costly and least valuable? Which requirements might the regulator be persuaded to drop? What proportion of R&D cost is a direct consequence of the 'throw money at it' tendency? In which therapeutic areas are molecular reductionism and brute force screening methods a distraction, and in which are they genuinely helpful? What explains the difference between these therapeutic areas? Perhaps the CDDO could quantify their analysis with a series of Venn diagrams like those in FIG. 3, to identify which sets and intersections have grown, and by how much, and which sets and intersections have shrunk. There should also be an attempt to measure the veracity of previous diagnostic and forecasting exercises. What has been the accuracy of internal forecasts on drug approvability and commercial success? Has this changed over time? What have been the most common kinds of error?

If the CDDOs provide a good explanation that is consistent with the idea that the countervailing forces will abate, or will be overcome, then all is well and good. If the explanation is unconvincing, or identifies forces that appear to be intractable, then the problems are obvious. At least it would advance the debate on how to balance the property rights of shareholders and the financial responsibilities of company boards with the wider benefits of safe, effective and affordable new drugs.

The prognosis for Eroom's Law

Just as we wanted to avoid proposing cures, we do not want to say too much about the prognosis for Eroom's Law. However, it might appear strange if we said nothing.

Despite the durability of the trend in FIG. 1, we would be surprised if Eroom's Law holds at an industry level over the next 5–7 years. Our view follows from two somewhat mechanical factors, in addition to one more interesting reason.

Turning to the first of the mechanical factors, the amount spent on R&D is not going to increase. The 'throw money at it' tendency is being tackled by most companies, with varying degrees of intensity. The second mechanical factor is the cumbersome biosimilar approval pathway that is emerging in the United States. Every aspect of the biosimilar production process can be scrutinized by the originator's lawyers, and this raises the prospect of endless blocking litigation. Consequently, developers of biosimilar products anticipate to get at least some of these products approved via the standard new biologics approval pathway (the FDA's biologics license application (BLA) process). These products will be approved as though they were novel agents, so they will inflate the number of novel approvals at very low R&D costs.

Turning to the interesting reason, we suspect that the signal-to-noise ratio may be improving among the compounds being developed for oncology indications. One or two other therapeutic areas may be similar in this respect. Perhaps there are hints of this in the FDA's new drug approvals in 2011. These totalled 30 overall, the most since 2004, although Munos[24] has shown that the distribution of new drugs approved by the FDA per year resembles the output of a Poisson process, so we do not want to over-interpret one good year (if new drug approvals did follow a Poisson process with a mean number of 26 from 1980 to 2010, we would expect 30 drugs to be approved by chance alone around once every 5 years). Looking in more depth at the nature of the 30 new drugs, eight were anticancer agents (brentuximab vedotin, vandetanib, crizotinib, ipilimumab, asparaginase, vemurafenib, ruxolitinib and abiraterone acetate). A focus on rare and poorly treated diseases is also visible in the 2011 total; 11 of the 30 new drugs were orphan drugs, and the orphan drugs included seven of the eight new anticancer agents. Orphan drugs are less prone to many of the factors discussed above, including the 'better than the Beatles' problem, the 'cautious regulator' problem and the big clinical trial problem.

Flat to declining R&D costs, as well as a bolus of oncology drugs, more orphan drugs and 'biosimilars as BLAs', might put an end to Eroom's Law at an industry level. Whether this improves things enough to provide decent financial returns on the industry's R&D investment is a different question. Financial markets don't think so. Industry executives do. It would be interesting to see what CDDOs think.

Jack W Scannell, Alex Blanckley and Helen Boldon are at Sanford C. Bernstein Limited, 50 Berkeley Street, Mayfair Place, London W1J 8SB, UK.

Brian Warrington is at Phoenix IP Ventures, 45 The Drive, Hertford, Hertfordshire SG14 3DE, UK.

Correspondence to J.W.S.
e-mail: Jack.Scannell@Bernstein.com

doi:10.1038/nrd3681

1. Hogan, J. C. Combinatorial chemistry in drug discovery. *Nature Biotech.* **15**, 328–330 (1997).
2. Geysen, H. M., Schoenen, F., Wagner, D. & Wagner, R. Combinatorial compound libraries for drug discovery: an ongoing challenge. *Nature Rev. Drug Discov.* **2**, 222–230 (2003).
3. [No authors listed.] Combinatorial chemistry. *Nature Biotech.* **18**, IT50–IT52 (2000).
4. Dolle, R. E. Historical overview of chemical library design. *Methods Mol. Biol.* **685**, 3–25 (2011).
5. Sanger, F. Sequences, sequences, and sequences. *Annu. Rev. Biochem.* **57**, 1–28 (1988).
6. Sanger, F. *et al.* Nucleotide sequence of bacteriophage phi X174 DNA. *Nature* **265**, 687–695 (1977).
7. Meldrum, C., Doyle, M. A. & Tothill, R. W. Next-generation sequencing for cancer diagnostics: a practical perspective. *Clin. Biochem. Rev.* **32**, 177–195 (2011).
8. Joachimiak, A. High-throughput crystallography for structural genomics. *Curr. Opin. Struct. Biol.* **19**, 573–584 (2009).
9. Van Brunt, J. Protein architecture: designing from the ground up. *Nature Biotech.* **4**, 277–283 (1986).

10. Mayr, L. M. & Fuerst, P. The future of high-throughput screening. *J. Biomol. Screen.* **13**, 443–448 (2008).
11. Schnee, J. E. Development cost: determinants and overruns. *J. Bus.* **45**, 347–374 (1972).
12. Baily, M. N. Research and development costs and returns: the U.S. pharmaceutical industry. *J. Polit. Econ.* **80**, 70–85 (1972).
13. Comanor, W. Research and technical change in the pharmaceutical industry. *Rev. Econ. Stat.* **47**, 182–190 (1965).
14. Grabowski, H. G., Vernon, J. M. & Thomas, L. G. Estimating the effects of regulation on innovation: an international comparative analysis of the pharmaceutical industry. *J. Law Econ.* **21**, 133–165 (1978).
15. Grabowski, H. & Vernon, J. A new look at the returns and risks to pharmaceutical R&D. *Manage. Sci.* **36**, 804–821 (1990).
16. Jensen, E. J. Research expenditures and the discovery of new drugs. *J. Ind. Econ.* **36**, 83–95 (1987).
17. Joglekar, P. & Paterson, M. L. A closer look at the returns and risks of pharmaceutical R&D. *J. Health Econ.* **5**, 153–177 (1986).
18. Elias, T., Gordian, M., Singh, N. & Zemmel, R. Why products fail in Phase III. *In Vivo* **24**, 49–56 (2006).
19. Pammolli, F., Magazzini, L. & Riccaboni, M. The productivity crisis in pharmaceutical R&D. *Nature Rev. Drug Discov.* **10**, 428–438 (2011).
20. Kola, I. & Landis, J. Can the pharmaceutical industry reduce attrition rates? *Nature Rev. Drug Discov.* **3**, 711–715 (2004).
21. DiMasi, J. A., Feldman, L., Seckler, A. & Wilson, A. Trends in risks associated with new drug development: success rates for investigational drugs. *Clin. Pharmacol. Ther.* **87**, 272–277 (2010).
22. Paul, S. M. *et al.* How to improve R&D productivity: the pharmaceutical industry's grand challenge. *Nature Rev. Drug Discov.* **9**, 203–214 (2010).
23. US Food and Drug Administration. Innovation or Stagnation: Challenge and Opportunity on the Critical Path to New Medical Products. *FDA website* [online], http://www.fda.gov/ScienceResearch/SpecialTopics/CriticalPathInitiative/CriticalPathOpportunitiesReports/ucm077262.htm (2004).
24. Munos, B. Lessons from 60 years of pharmaceutical innovation. *Nature Rev. Drug Discov.* **8**, 959–968 (2010).
25. Borhani, D. W. & Butts, J. A. Rethinking clinical trials: biology's mysteries. *Science* **334**, 1346–1347 (2011).
26. David, E., Tramontin, T. & Zemmel, R. Pharmaceutical R&D: the road to positive returns. *Nature Rev. Drug Discov.* **8**, 609–610 (2009).
27. Garnier, J. P. Rebuilding the R&D engine in big pharma. *Harv. Bus. Rev.* **86**, 68–79 (2008).
28. Agarwal, S. *et al.* Unlocking the value in big pharma. *McKinsey Quarterly* **2**, 65–73 (2001).
29. Ruffolo, R. R. Engineering success: Wyeth redefines its research & development organisation. *Drug Discovery World website* [online], http://www.ddw-online.com/s/business/p148328/engineering-sucess:-wyeth-redefines-its-research-&-development-organisation-fall-05.html (2005).
30. Douglas, F. L., Narayanan, V. K., Mitchell, L. & Litan, R. E. The case for entrepreneurship in R&D in the pharmaceutical industry. *Nature Rev. Drug Discov.* **9**, 683–689 (2010).
31. Zhong, X. & Moseley, G. B. Mission possible: managing innovation in drug discovery. *Nature Biotech.* **25**, 945–946 (2007).
32. Horrobin, D. Realism in drug discovery — could Cassandra be right? *Nature Biotech.* **19**, 1099–1100 (2001).
33. Horrobin, D. F. Innovation in the pharmaceutical industry. *J. R. Soc. Med.* **93**, 341–345 (2000).
34. Horrobin, D. F. Modern biomedical research: an internally self-consistent universe with little contact with medical reality? *Nature Rev. Drug Discov.* **2**, 151–154 (2003).
35. Ruffolo, R. R. Why has R&D productivity declined in the pharmaceutical industry? *Expert Opin. Drug Discov.* **1** 99–102 (2006).
36. Le Fanu, J. *The Rise and Fall of Modern Medicine* (Little Brown, London, 1999).
37. Pisano, G. *Science Business: The Promise, the Reality, and the Future of Biotech.* (Harvard Business School Press, Boston, 2006).
38. Young, M. P. Prediction v Attrition. *Drug Discovery World website* [online], http://www.ddw-online.com/s/business/p92811/prediction-v-attrition-fall-08.html (2008).
39. Hopkins, A. L., Mason, J. S. & Overington, J. P. Can we rationally design promiscuous drugs? *Curr. Opin. Struct. Biol.* **16**, 127–136 (2006).
40. Tollman, P., Morieux, Y., Murphy, J. K. & Schulze, U. Identifying R&D outliers. *Nature Rev. Drug Discov.* **10**, 653–654 (2011).
41. Ford, E. S. *et al.* Explaining the decrease in U.S. deaths from coronary disease, 1980–2000. *N. Engl. J. Med.* **356**, 2388–2398 (2007).
42. Lichtenberg, F. The impact of drug launches on longevity: evidence from longitudinal disease-level data from 52 countries, 1982–2001. *Int. J. Health Care Finance Econ.* **5**, 47–73 (2005).
43. Schnee, J. E. R&D strategy in the U.S. pharmaceutical industry. *Res. Policy* **8**, 364–382 (1979).
44. Hopkins, A. L. & Groom, C. R. The druggable genome. *Nature Rev. Drug Discov.* **1**, 727–730 (2002).
45. Russ, A. P. & Lampel, S. The druggable genome: an update. *Drug Discov. Today* **10**, 1607–1610 (2005).
46. Overington, J. P., Al-Lazikani, B. & Hopkins, A. L. How many drug targets are there? *Nature Rev. Drug Discov.* **5**, 993–996 (2006).
47. Roth, B. L., Sheffer, D. L. & Kroeze, W. K. Magic shotguns versus magic bullets: selectively non-selective drugs for mood disorders and schizophrenia. *Nature Rev. Drug Discov.* **3**, 353–359 (2004).
48. Wurtman, R. J. & Bettiker, R. L. The slowing of treatment discovery, 1965–1995. *Nature Med.* **1**, 1122–1125 (1995).
49. Healy, D. *The Psychopharmacologists: Volume 2* 93–118 (Hodder Arnold, London, 1999).
50. Healy, D. *The Psychopharmacologists: Volume 2* 259–264 (Hodder Arnold, London, 1999).
51. Healy, D. *The Antidepressant Era* (Harvard University Press, Cambridge, Massachusetts, 1997).
52. Weatherall, M. An end to the search for new drugs? *Nature* **296**, 387–390 (1982).
53. Richard, J. & Wurtman, M. D. What went right: why is HIV a treatable infection? *Nature Med.* **3**, 714–717 (1997).
54. [No authors listed.] A dearth of new drugs. *Nature* **283**, 609 (1980).
55. Persson, C. G., Erjefält, J. S., Uller, L., Andersson, M. & Greiff, L. Unbalanced research. *Trends Pharmacol. Sci.* **22**, 538–541 (2001).
56. Ainsworth, C. Networking for new drugs. *Nature Med.* **17**, 1166–1168 (2011).
57. Denome, S. A., Elf, P. K., Henderson, T. A., Nelson, D. E. & Young, K. D. *Escherichia coli* mutants lacking all possible combinations of eight penicillin binding proteins: viability, characteristics, and implications for peptidoglycan synthesis. *J. Bacteriol.* **181**, 3981–3993 (1999).
58. Keith, C. T., Borisy, A. A. & Stockwell, B. R. Multicomponent therapeutics for networked systems. *Nature Rev. Drug Discov.* **4**, 71–78 (2005).
59. Lombardino, J. G. & Lowe, J. A. The role of the medicinal chemist in drug discovery — then and now. *Nature Rev. Drug Discov.* **3**, 853–862 (2004).
60. Swinney, D. C. & Anthony, J. How were new medicines discovered? *Nature Rev. Drug Discov.* **10**, 507–519 (2011).
61. Reichert, J. M. Probabilities of success for antibody therapeutics. *mAbs* **1**, 387–389 (2009).
62. Steward, F. & Wibberly, G. Drug innovation — what's slowing it down? *Nature* **284**, 118–120 (1980).
63. Collins, F. S. Medical and societal consequences of the Human Genome Project. *N. Engl. J. Med.* **341**, 28–37 (1999).
64. Rees, J. Post-genome integrative biology: so that's what they call clinical science. *Clin. Med.* **1**, 393–400 (2001).
65. Grove, A. Rethinking clinical trials. *Science* **333**, 1679 (2011).
66. O'Shaughnessy, J. *et al.* Iniparib plus chemotherapy in metastatic triple-negative breast cancer. *N. Engl. J. Med.* **364**, 205–214 (2011).
67. O'Shaughnessy, J. *et al.* A randomized Phase III study of iniparib (BSI-201) in combination with gemcitabine/carboplatin (G/C) in metastatic triple-negative breast cancer (TNBC). *J. Clin. Oncol.* **29**, Abstr. 1007 (2011).
68. Guha, M. PARP inhibitors stumble in breast cancer. *Nature Biotech.* **29**, 373–374 (2011).
69. Soda, M. *et al.* Identification of the transforming *EML4–ALK* fusion gene in non-small-cell lung cancer. *Nature* **448**, 561–566 (2007).
70. Chapman, P. B. *et al.* Improved survival with vemurafenib in melanoma with BRAF V600E mutation. *N. Engl. J. Med.* **364**, 2507–2516 (2011).
71. [No authors listed.] Regulatory watch: leading hedgehog inhibitor submitted for approval as skin cancer drug. *Nature Rev. Drug Discov.* **10**, 802–803 (2011).
72. DeMonaco, H. J., Ali, A. & von Hippel, E. The major role of clinicians in the discovery of off-label drug therapies. *Pharmacotherapy* **26**, 323–332 (2006).
73. Mathieu, M. P. (ed.) *Parexel's Bio/Pharmaceutical R&D Statistical Sourcebook 2010/2011* 163–261 (Barnett International, Needham, Massachusetts, 2010).
74. Marshall, G. *et al.* Streptomycin treatment of pulmonary tuberculosis. *BMJ* **30**, 769–782 (1948).
75. MacNeil, J. S. H. *Changes in the characteristics of approved New Drug Applications for antihypertensives.* Thesis, Massachusetts Institute of Technology (2007).
76. Lin, H. S. *Changes in the characteristics of new drug applications for the treatment and prevention of diabetes mellitus.* Thesis, Massachusetts Institute of Technology (2007).
77. Scandinavian Simvastatin Survival Study Group. Randomised trial of cholesterol lowering in 4444 patients with coronary heart disease: the Scandinavian simvastatin survival study (4S). *Lancet* **344**, 1383–1389 (1994).
78. Munos, B. How to avert biopharma's R&D crisis. *In Vivo* **29**, 2011800050 (2011).
79. Macarron, R. *et al.* Impact of high-throughput screening in biomedical research. *Nature Rev. Drug Discov.* **10**, 188–195 (2011).
80. Bohacek, R. S., McMartin, C. & Guida, W. C. The art and practice of structure-based drug design: a molecular modeling perspective. *Med. Res. Rev.* **16**, 3–50 (1996).
81. Brown, D. Future pathways for combinatorial chemistry. *Mol. Divers.* **2**, 217–222 (1996).
82. Lipinski, C. A., Lombardo, F., Dominy, B. W. & Feeney, P. J. Experimental and computational approaches to estimate solubility and permeability in drug discovery and development settings. *Adv. Drug Deliv. Rev.* **46**, 3–26 (2001).
83. Gleeson, M. P., Hersey, A., Montanari, D. & Overington, J. Probing the links between *in vitro* potency, ADMET and physicochemical parameters. *Nature Rev. Drug Discov.* **10**, 197–208 (2011).
84. Kay, J. *Obliquity: Why our goals are best achieved indirectly* (Profile Books, London, 2010).
85. Watts, D. J. & Strogatz, S. H. Collective dynamics of 'small-world' networks. *Nature* **393**, 440–442 (1998).
86. Pharmaceutical Research and Manufacturers of America. Pharmaceutical Industry Profile 2011. *PhRMA website* [online], http://www.phrma.org/sites/default/files/159/phrma_profile_2011_final.pdf (Washington DC, PhRMA, April 2011).
87. Congress of the United States: Congressional Budget Office. Research and Development in the Pharmaceutical Industry. *Congressional Budget Office (CBO) website* [online], http://www.cbo.gov/ftpdocs/76xx/doc7615/10-02-DrugR-D.pdf (October 2006).

Acknowledgements

W. Bains, T. Curtis, B. Charlton, M. Young, O. Imasogie, G. Porges, and B. Munos were generous with their time and ideas during various stages in the genesis of this article.

Disclaimer

The information provided herein was prepared by Sanford C. Bernstein & Co. LLC and Brian Warrington. It is not investment research, although it may refer to a Bernstein research report or the views of a Bernstein research analyst. This communication does not constitute a complete fundamental analysis of any companies mentioned. Unless indicated, all views expressed herein are the views of the authors and may differ from or conflict with those of the Bernstein Research Department. This article does not constitute investment advice or recommendations and is not a solicitation or an offer to purchase or sell securities. The information contained herein is only good as of the date and time of the publication; we do not undertake to advise of any changes in the opinions or information contained herein.

Competing interests statement

The authors declare competing financial interests: see Web version for details.

FURTHER INFORMATION

RCSB Protein Data Bank database:
http://www.rcsb.org/pdb/statistics/holdings.do

SUPPLEMENTARY INFORMATION

See online article: S1 (table) | S2 (box)

ALL LINKS ARE ACTIVE IN THE ONLINE PDF