The Outcomes of 25 years of Green Chemistry.

Steven Abbott
steven@stevenabbott.co.uk

Background

A core tenet of sustainability is that resources are limited and so should be used wisely to create optimal outcomes. After 25 years of Green Chemistry, a prestigious journal in the green community, we can look at the titles and abstracts of 11,000 published papers to try to objectively assess the outcomes of the resources devoted to the work in the journal. Specifically, how well has it achieved the implicit goals of creating chemistry, processes and products that are "good for the planet"? The evidence presented in this personal view of the outcomes suggests that the answer is not encouraging at any meaningful scale. Because the data and tools used in my own analysis are provided in a Creative Commons, Open Access, Excel document, others can use that resource to reach their own conclusions and publish them as they wish. Indeed, views opposing my generally pessimistic view will be especially welcomed.

Although even a flagship journal such as Green Chemistry cannot capture all the efforts within the green chemistry community, it seems reasonable to assume that it represents many of the key themes and outcomes of the larger community. Feel free to disagree.

Methodology

On 3 December 2024, the full records (including Titles and Abstracts) of all 11,500 papers defined in Web of Science as being in Green Chemistry since its inception in 1999 were downloaded in Excel format. This is an arbitrary cutoff date but captures the essence of ~25 years of the journal.

As mentioned above, all data and analyses are available in an Excel file that can be downloaded along with this paper.

The dataset then had to be curated down to the 11,000 used in the analysis which follows. 17 papers wrongly ascribed to the journal (from 1996) were removed, 262 with page numbers marked as "T" or "G" were removed as these were more commentary than research. 203 were removed either because they were single page or were corrections. Finally, the most recent 83 papers were removed as they were advanced publications lacking page numbers. "Removal" in this instance means moving to separate worksheets in the Excel workbook for those who wish to mine them for extra information.

The citation count for each paper was part of the download. Extra columns were added using standard Excel formulae to calculate number of authors and number of pages per paper, calculated from the provided Start and End pages. In order to make the page count metric complete, start/end values were found directly from the RSC site for about 100 papers which, for whatever reason, lacked one or both of these values.

Basic data could be calculated as shown in the Table:

	Number of Authors	Number of Pages	Citation Count
Mean	5.5	9.4	62.5
Median	5	8	34
Mode	4	6	0

Then histograms of Number of Authors, Page Count, Citations and 180 Day Usage were created. The graphic gives a visual summary; some details are discussed later. Histogram views are necessarily compromises and readers are encouraged to go into the Excel file and change the axes to produce views that address their own concerns. For example, a tweak to the x-axis will let you see how many papers from the 2000+ with 0 to 10 citations had fewer than 5, or how many from the more than 4000 with 0 or 1 180 day usage had none.

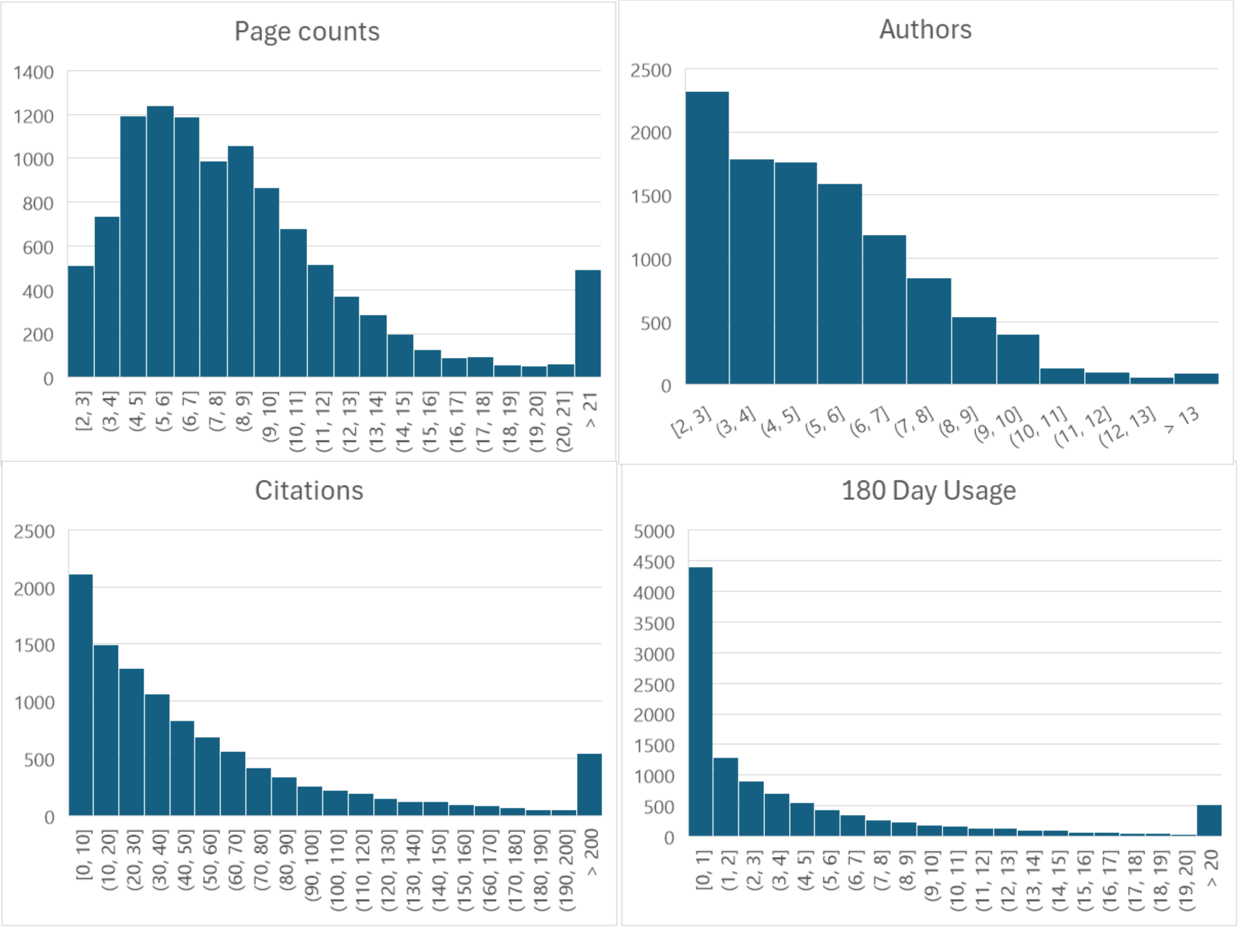

Extra plots can be made in the Excel file. For example, a graph of number of citations versus number of authors (not shown) indicates no correlation, with some highly cited papers having few authors and some with low citation count having many authors.

A key concern of Green Chemistry is global CO_2 emissions, so I regret that I cannot find a way to estimate the reduction in CO_2 enabled by the published science. My own work (see below) on typical bio-based products shows that savings will be negligible, but this is a narrow view of a broad field. It would be wonderful if some group with the required expertise could provide some figures. But arguably the green publication community should be doing these calculations routinely so we could readily assess the potential outcomes of their work.

The work itself created CO_2 which should be accounted for. Using a published estimated figure (André Estevez-Torres et al, *Carbon footprint and mitigation strategies of three chemistry laboratories*, Green Chem., 2024, 26, 2613-2622) of 5 tCO_2pa per lab worker, and making the assumption that lab headcount was equivalent to 50% of the number of authors, and that each paper represented 1 year of work, we arrive at 150 $ktCO_2$ emitted by this research over 25 years. This figure depends on multiple assumptions which the reader can alter as they wish. As we shall see, $ktCO_2$ are irrelevant to our planet-scale concerns, so the impact of the CO_2 generated by the work itself has been negligible.

Similarly, each of the 50% of the authors working for 1 year per paper amounts to a headcount cost of 1 full-time-equivalent year; so assuming $50,000 (constant 2024 $) per full time equivalent per year, over the 11,000 papers and 30,000 person years, the expenditure is of the order of $1.5 billion. Again, that figure depends on assumptions, but it is hard to find reasons to reduce this to below $1 billion. So we need to ask whether the resources of 30,000 person years and $1+ billion have resulted in significant green outcomes.

To explore key themes, single and bigram word counts of all titles and all abstracts were obtained using public domain websites. These are provided in the spreadsheet. It is easy to look beyond the trivial words such as “the” to get a good feel for the concerns of Green Chemistry authors.

Using those words/phrases as inspiration it was possible to map such themes during the journal’s 25 years. A phrase such as “ionic liquid” (which, of course, will also detect “ionic liquids”) or a check for two versions of the same phrase, e.g. “biobased” and “bio-based” (double counts are excluded) can be entered and a graph created. Because the number of papers per year has increased steadily from ~100 in early years to 800+ in recent years, the graphs can be plotted as a count per year or, for a fairer comparison, as a normalised value, as shown here:

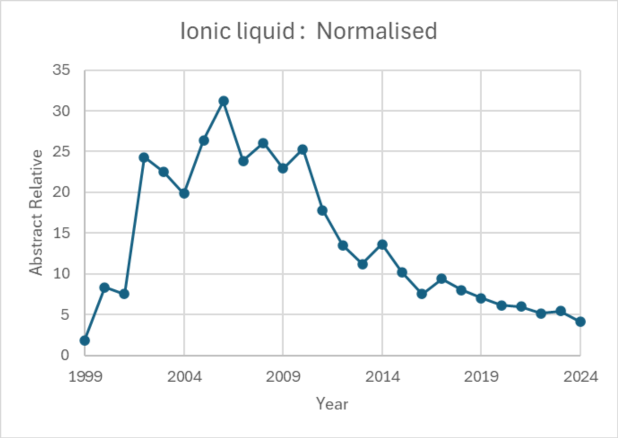

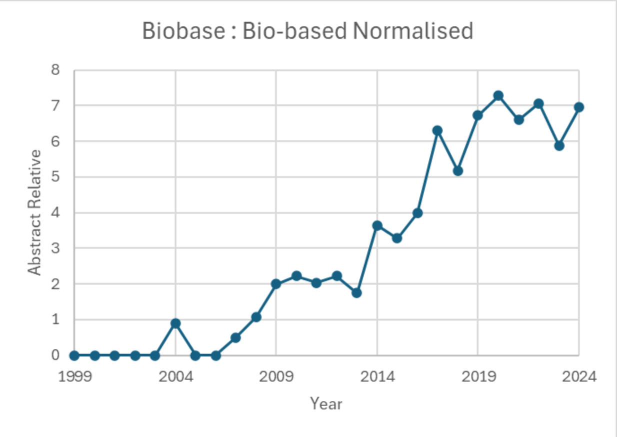

Users can also enter a number of phrases and see their total counts (Titles and Abstracts) over the years. For convenience with a complex Excel formula, ### is used to indicate only 1 search term. This is a snapshot of terms I happened to be using at the time of writing:

Terms	Alternative	In title	In abstract
Ionic liquid	###	1004	1170
deep eutectic	eutectic solvent	194	254
Supercritical	###	168	249
Biobased	Bio-based	235	498
Polylactic acid	Poly(lactic acid)	14	29
Biopolymer	Bio-polymer	31	158
Biosurfactant	###	4	16
Green solvent	###	32	250
Sustainab	###	674	1912
Life cycle	Life-cycle	90	230
CO_2	carbon dioxide	869	1326
Toxic	###	105	850
Cataly	###	4099	6315
Enzym	###	288	789
Carbon diox	###	235	362
MOF	metal-organic	103	149
Recycl	###	429	1524
circular	###	21	132
carbon dioxide	###	235	362
CO_2	Supercritic	769	1289
pyrrolidone	NMP	18	64
one pot	one-pot	280	668

A note on AI

I used ChatGPT only once – to create an Excel formula that I couldn't work out on my own. Those with greater AI skills will surely be able to analyse the data with more insight than simple word counts can manage. Again, the Excel file is open source, and I would welcome alternative, more sophisticated, analyses.

Outcomes (Titles/Abstracts)

We can now look at those outcomes from the research about which I have some knowledge, using key phrases followed by the counts of their appearances in Title/Abstract

Outcomes – Ionic Liquids (1004/1170)

Almost 10% of the papers (and the most highly-cited paper) were about ionic liquids. A wonderful 2016 review (Werner Kunz & Katharina Häckl, *The hype with ionic liquids as solvents*, Chemical Physics Letters 661 (2016) 6–12) sums up a lack of critical analysis by the green chemistry community around these solvents called "green" largely because of their low volatility. A 2020 review (Adam J. Greer et al, *Industrial Applications of Ionic Liquids*, Molecules 2020, 25, 5207) pointed out that large-scale uses of ILs as solvents has to date not materialised. As Kunz emphasises, ILs show valuable properties for specialist applications, some of which are described in the Greer review, but they could never, and have never, become useful green solvents, at scale, to replace non-green solvents. One interpretation of the graph above might be that this was largely recognised within the green chemistry community after 2010. One question might be if the problems with ILs could have been spotted earlier?

A graph (not shown) of Deep Eutectic Solvents shows enthusiasm starting around the decline of interest in ILs and continuing to the present. And clearly, as with ILs, using them when they provide unique advantages is nice, but caution should be applied when considering how much impact on a planetary scale these might ever be expected to have. The IL story suggests that deep eutectics, with problems such as high viscosity or alarming skin safety for keratin-dissolving formulations, might be more of a bandwagon than a serious attempt to solve problems of planetary significance.

Outcomes – Solvent substitutions (32/250)

Solvent substitution has been a pervasive theme of Green Chemistry. Who doesn't want zero VOC solvents, easily recycled via evaporation/condensation, safe, bio-based, bio-degradable, sustainable, low odour and low cost? On the other hand, however green a solvent might be, it is useless if it cannot dissolve the required solute, or if it is too expensive, or not available at scale. Many of the simpler greener solvents (such as FAME) are in parts of solubility space not much needed by industry and/or are large molecules that are thermodynamically less effective. This is an especially acute problem when it comes to replacing the dipolar aprotic solvents which are required for many processes.

A 2021 review (Neil Winterton, *The green solvent: a critical perspective*, Clean Technologies and Environmental Policy (2021) 23:2499–2522) points out that the inherent contradictions mean that the number of meaningful solvent substitutions *at scale* has been negligible.

The story of Cyrene as a green alternative to NMP and other dipolar aprotic solvents highlights the huge risks involved in taking a promising bio-based solvent, with a desirable toxicological profile, and scaling it up. A single failure of a toxicity test (which, happily, did not happen) could have rendered useless the €millions spent on REACH certification. In the end, some quirks of the molecule, the costs of upscaling, and the lack of a single large market meant that the considerable scientific, regulatory, engineering and financial investments have to date not yielded a secure supply on the market, and the company involved has filed for bankruptcy. Others who aspire to introduce novel green solvents at scale will be discouraged by this case study.

A relatively little-known outcome (see the discussion in my free online Solubility Science book, https://www.stevenabbott.co.uk/practical-solubility/the-book.php) of an EU project on bio-based butyl acetate as a replacement for the current large-scale petrochemical-based material, was the Life Cycle Analysis (LCA) showing that benefits/drawbacks were mixed. The LCA showed that it was not at all obvious that the planet would be better-off with the bio-based version. The use of food-crop ethanol to make a chemical that can be readily produced at high efficiency in a standard petrochemical plant is not without its detractors.

An interesting observation by Winterton is relevant. Much of green chemistry research looks at solvents that might be greener for the *synthesis* step of a molecule – ignoring the fact that volumes of solvent for the unglamorous "workup" procedures are generally more significant, and issues such as easy recovery of the solvent via distillation are proportionately more important for manufacturers. The impact of a greener reaction solvent on the overall solvent use would be less significant in terms of the planet.

Even water, as the ultimate green solvent has the well-known problem of a 4x higher enthalpy of vaporisation. Again, Winterton points out that manufacturers are keenly aware of the cost of energy, and expending 4x more to evaporate a solvent is not welcomed.

The push to replace trichloroethylene pre-dates Green Chemistry and has largely succeeded. Replacing dichloromethane is a continuing challenge. It is an excellent, efficient solvent with a long history of safe use in responsible industrial environments. Finding a direct replacement is impossible. It is a tough challenge to find a combination of greener solvents and procedures that take into account slower solvation action and difficulty in recycling via distillation.

Propylene carbonate is an interesting green solvent. Its low vapour pressure defines it as zero VOC in many jurisdictions and the fact that it can, in principle, be derived from CO_2 gives an aura of carbon neutrality (even negativity). However, it should be noted that the current use of ~0.5 Mt of propylene carbonate would only store a few minutes' worth of CO_2. And currently it is produced with high efficiency via phosgene, not via direct CO_2 capture. Its low vapour pressure makes used propylene carbonate difficult to recycle via distillation. At end of life, it easily reverts to propylene glycol and CO_2. For those who *need* a low vapour pressure solvent in an interesting region of solubility space (e.g. in batteries), it is an easy choice.

One paper in Green Chemistry (Denis Prat et al, *CHEM21 selection guide of classical- and less classical-solvents*, Green Chem., 2016, 18, 288) stands out for its alternative approach (and for its citation count). The Chem21 consortium of industry and academia took the view that industry should be able to choose "least bad" solvents that do what is required of them. So across a range of solvation problems common to industry, involving solutes from least polar to most polar, from aliphatic to aromatic, aprotic to protic, solvents were chosen for their, safety, green credentials and more. No one much likes using acetonitrile in large quantities, but among the dipolar aprotic solvents it was judged to be the least problematic. The Chem21 list – which is broadly aligned with many other specific guides produced by individual companies – allows those in industry (i.e. most users of solvents) who want to mix efficacy with green credentials to use solvents judged by experts to be acceptable in the real world of necessary compromises.

This paper does not have lofty planet saving goals, but simply aims to provide help to those in industry who have to make difficult choices in their move to greener working.

Outcomes – scCO_2 (174/249)

Supercritical CO_2 on its own is well-recognised as a poor solvent, on a par with i-pentane. So the approach of heating/pressurising the CO_2, extracting, then releasing the extract by releasing the pressure, faces significant challenges / restrictions to scope. Instead the process generally requires "entrainers", molecules at a few % concentration that can render target molecules more soluble, with the downside of the need to remove the entrainer molecule for safety or economic (recycling) reasons.

A review of the field (Steven Abbott & Seishi Shimizu, *Understanding Entrainer Effects in Suprcritical CO_2*, Ch 2 in Andrew J Hunt & Thomas M Attard, *Supercritical and Other High-pressure Solvent Systems*, 2018, Royal Society of Chemistry) revealed a cavalier approach to efficiency in thinking through whether scCO_2 was the best choice and, if so, what was the best entrainer.

Caffeine *can* be removed from coffee using scCO_2, but it requires water as the entrainer. Around the world, ethyl acetate (called the sugar method, because the solvent is prepared from fermentation of sugary feedstocks – an interesting marketing "spin"), dichloromethane, liquid CO_2 and water ("Swiss water process") are much-used alternatives because decaffeination has many issues other than which is the greenest solvent. Given the large scale of each of the 4 alternatives, it's not obvious that scCO_2 has been a great success.

Although there are plenty of small-scale uses of $scCO_2$, it's unclear whether these uses are because the approach is green, or just the best balance of extraction efficiency with affordability at the specific scale. I can find no evidence of impact at scale. The boom and bust of $scCO_2$ for extraction of cannabinoids is a cautionary tale.

In terms of impact on the planet, the $scCO_2$ review arguably shows a misalignment between current incentives in science and our global goals. Papers always started by saying how green $scCO_2$ is, thereby justifying the production of a paper along with their specific experiments. Instead of a relatively few experiments that would systematically prove/refute an attractive hypothesis from the 1990's on how entrainers work, each paper chose a unique (ultimately flawed) theory for why $scCO_2$+entrainer worked and carried out uncoordinated experiments which did not create a reliable corpus of evidence. The statistical thermodynamic approach of Shimizu showed that the 1990's hypothesis was correct, explained why it was correct and provided some simple rules to enable wise choice of entrainer.

The $scCO_2$ story is another example of confusion of good intentions with good outcomes. A proper determination to be green could have enabled the $scCO_2$ community to understand the basic rules decades earlier, avoiding much waste of precious resource. The route to small-scale, specialist applications would have been greener.

Outcomes – Bio-based feedstocks (235/498), Lignin (566/850)

Companies can measure the C13:C12 ratio of a product to indicate the % bio-based origin. Unfortunately, such measures cannot say how much of that C13 came from food crop resources rather than "waste". Given the ethical concerns over using food crops as feedstock, just being bio-based is not good enough, so precious resource needs to go in to proving and policing (fraud is ubiquitous) the integrity of the supply chain.

Additionally, agrochemical "waste" is not as available as naïve hopes suggested – the majority gets converted to some use and/or is returned to keep carbon in the soil. If the wastes have to be aggregated to large conversion facilities to make chemicals at scale, their bulk (low-density) makes that aggregation problematical. And there are arguments that the energy transition could be assisted by waste (including food waste) converted to bio-gas by efficient digestion systems.

Popular wastes such as "used cooking oil" must be put into context. For Sustainable Air Fuels, SAF: "The actual share of aviation needs that can be met by these [waste oil and fats] is almost certainly less than 2%, even if these raw materials are entirely used for this purpose, rather than existing uses." (https://www.carboncommentary.com/blog/2023/12/1/waste-oils-will-not-provide-substantial-volumes-of-sustainable-aviation-fuel-despite-what-mr-sunak-says).

And the common expression used about the even more popular waste, lignin, is "You can make everything from lignin except money". This isn't through want of trying as the numbers in the dataset suggest. The effort seems to continue unabated despite near-universal discouraging results with this intractable material:

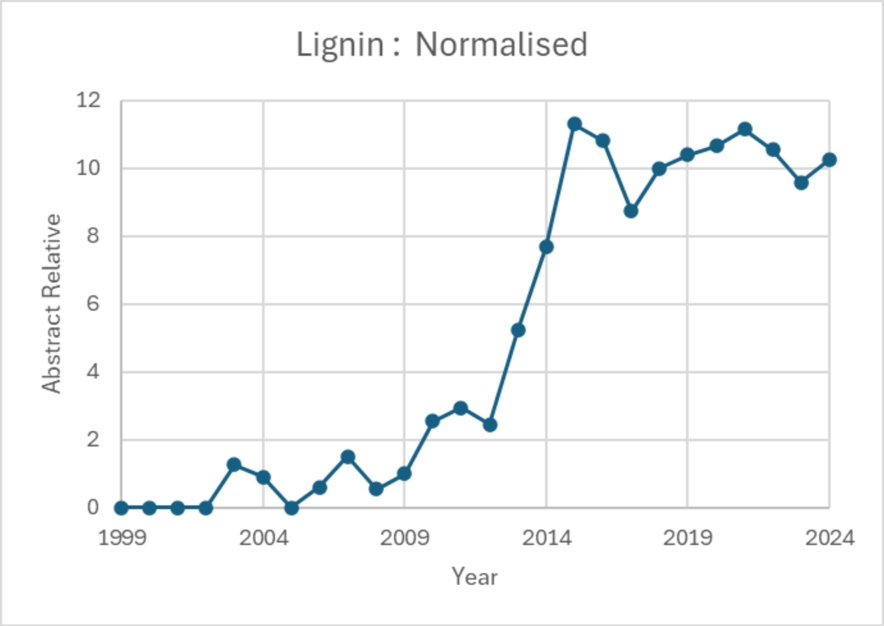

Outcomes – Bio(-based) polymers (31/158)

It has long since been obvious to the green chemistry community that bio-based polymers or biopolymers produced directly from fermentation, or polymers created (somehow) from CO_2 *must* be superior to petrochemical polymers.

A few moments of reflection might have saved a lot of resource wasted on hopeless quests.

The reason we use (mostly) PE, PP, PET, PC, PS, ABS, PUs and polyamides is that decades of research into alternatives from petrochemical resources failed to find polymers with the price/performance ratio of these select few. It is survival of the fittest based on their fitness for different purposes. The chances that any bio- or bio-based polymer will approach any of these in potential utility are small, because the chances of *any* alternative polymer being competitive are provably small.

One challenge comes when marketing oversells a polymers green credentials, leading to problems for the whole field. There has been an oversimplified claim that a polymer such as polylactic acid, PLA, is intrinsically green because it:

- can come straight from fermentation;
- is an adequate polyester;
- can revert to lactic acid if thrown away as trash, or sent to a municipal composting facility.

The first problem with PLA is that it is potentially produced using food crops. Next, “adequate” is not good enough. The famous Frito Lay debacle where consumers hated the “crinkly” sound of crisp/chip packets is a classic example of the challenges faced when bringing new products to market. It is well-recognised that its popularity in 3D printing is its convenience rather than the mechanical and thermal stability of the printed parts. Finally, its “compostability” was theoretical – in real-world systems it often leaves large residues. In any case, consumers were so confident of its biodegradability that they would happily throw it as open trash – where it provably does not fall apart. Although this has been known for some years, on the day of editing this section, the UK Advertising Agency banned an advert for a PLA coffee pod which were claimed to be “compostable”. It ruled that consumers would be misled into thinking they could compost the pods in their garden or their municipal composting facility.

PLA’s specific uses in medical applications are based on its considerable merits around tunable degradability – such as sutures that disappear after the desired number of weeks. However, I would argue this isn’t so much “green chemistry” as “smart use of specific properties because they provide a core functionality”.

We come to recycling shortly, but one way to improve polymer recycling is to have fewer polymers to be sorted, reducing the issues of cross-contamination. PLA, PBS and the other “green polymers” are, from this viewpoint, happy failures as their quantities are too small to significantly add to the problems. Imagine the problem if PEF (polyethylene furanoate) became available in large quantities. It is, indeed, an excellent replacement for PET, but we would now have the problem of sorting PEF bottles from PET bottles, with inevitable cross-contamination in the resulting recycled polymers. The issue is that those promoting bio-based polymers failed to think through the fact that their “solution” would itself creating problems.

In terms of emotive topics such as the Pacific Garbage Patch, no green polymer proposal will have any impact in any timescale that makes any sense. The one investment, at scale, which will reduce the problem is unglamorous garbage collection in those many countries that have poor or zero garbage collection infrastructure. Collection rather than throwing into the nearest river (through lack of alternatives) is the root-cause solution. What to do with that garbage is part of the next topic.

Finally we have the classic green-washing around bio-based polymers: the famous “plant bottle” for fizzy drinks. Taking sugar from a food source, converting it to ethanol, converting that to ethylene glycol then combining that with petrochemical terephthalic acid (there is an obvious lack of useful bio-based aromatics) never made sense other than to marketeers. Arguably the green chemistry community did a poor job of resisting this – after all, the assumption behind many published papers is that if it’s a bio-based polymer it *must* be good.

Outcomes – Recycling (429/1524)

It is one thing to claim that “X is recyclable”. It is another to fit X into the present or likely future realities of recycling. If X can be recycled via process A, Y by process B and Z by process C, each of these recyclable materials is adding to an already over-complex system of ill-thought-through recycling schemes. It is as though “recyclable” is a magic word – invoke it and the chemical or polymer will, indeed, be recycled.

What is distressing is the lack of serious attention by the green chemistry community to the issue of obtaining real outcomes rather than producing feel-good papers.

A relatively low-hanging fruit could be argued to be PET bottles. They are relatively clean, easily identified because they exist in standard form factors, and a few countries (e.g. Switzerland) had the foresight to ensure that they were separated and collected at scale. Because resources are precious, the green chemistry community could have engaged in a lively scientific debate on the best way to process the returned plastic. Downcycle, eat it with enzymes, use strong caustic, use ionic liquids – they each have upsides and downsides. But even for PET we currently have no scientific consensus on the optimal way forward, so it struggles to be recycled at scale with the best-available technology. As mentioned above, well-meaning attempts to introduce PEF would further complicate matters.

Seriously focussing on PET from bottles would help with the trickier problem of recycling vast quantities of polyester fibres. In the short term, the solution is to discourage "fast fashion", an approach outside the skillset of green chemistry, perhaps explaining the low count of "textiles" (15/38) in the dataset. The fact that much polyester is mixed with cotton is an extra challenge. In this context, the natural, rather intractable cotton makes a tough challenge tougher.

PE and PP make:

- superb, lightweight bottles and jars for food and personal care products;
- barrier packaging films for food and pharma;
- larger structural components for aviation and automobiles where high-quality material is crucial for light-weighting.

They are not going to be replaced in any relevant timescale. And downcycled product is of little or dubious use. Advanced Chemical Recycling (ACR) can readily rip the polymers apart to an "oil" which can be fed back into efficient crackers to create virgin polymers. There is an energy/CO_2 cost for doing so, but there are no alternatives at scale.

What makes ACR hard is that mixes of PE/PP with other polymers become more difficult to reliably convert into usable oil. An intensive effort to improve the quality of the feedstock via rational waste sorting, plus intensive research into catalyst systems able to cope with impurities in the feed, would allow recycling at scale.

Instead, the green chemistry community takes it for granted that PE/PP are "non-recyclable" and insist on attempting to introduce more complexity into the system with polymers that entirely lack the proven capabilities and massive existing infrastructure behind PE/PP. Instead of solving a problem at scale, the community tries to introduce more polymers into the recycling mix, making the problem even harder to solve.

One irony in the recycling world is the requirement for personal care companies to use PCR – Post Consumer Recycled – plastic. It is quite a shock to see examples of, say, PCR shampoo bottles. There is now a whole sub-industry working on the potential risks of using non-virgin grades in consumer products. For those using PE, a great source of relatively high-grade PCR PE is UK milk bottles. These are, of course, virgin PE and are available in large quantities that are relatively easy to clean and recycle (though only relatively). At the time of writing, in the EU personal care companies cannot use virgin grade polymer produced via ACR because it is not officially classed as a form of PCR.

In another irony, a large Swiss chocolate company is switching over their chocolate bars to paper wrappers instead of plastic. We all know that "paper is recyclable". But in Switzerland, 100% of those paper wrappers will end up being incinerated for fuel, just as currently happens to the plastic wrappers.

You can see the good intentions behind PCR and paper wrappers, but the good outcomes are not so obvious.

Outcomes – Renewability (222/1005)

To say that something is "renewable" can bestow an immediate green credential. More recently the term has had to be nuanced to reflect that renewable resources that compete with food crops are undesirable in many ways. And we have already looked at the problems of "waste" materials. It is hard to find in the dataset wonderful examples of the use at scale of truly renewable resources.

What about a 100% natural resource that is quasi-renewable or substantially recyclable, at scale, using existing infrastructure? We are never going to run out of all-natural oil from underground sources *once we stop burning it as fuel*. The entire chemical industry uses less than 2% of annual world oil consumption. So if, at current oil consumption rates, we were going to run out in, say, 10 years (a ridiculous number), and if we used it instead for chemicals with zero recycling (which would be stupid), it would take us 500 years to use it up.

Suppose, therefore, that the green chemistry community devoted its intellectual fire power to the problems of, at first, recycling relatively clean industrial waste PE/PP with, perhaps, extra feeds of well-sorted post-consumer plastics. There are numerous challenges of efficiency, catalysis, CO_2 emissions, handling of intractable residues, which the chemistry community is well-equipped to solve. Then the bigger challenges of handling badly-sorted post-consumer waste can be faced; this will require highly sophisticated catalyst systems and smart handling of the (mostly) oxygenated contaminants.

There is a natural bias against "big petrochemical" because these corporations have often been part of the problem and less often are part of the solution. But the environment needs solutions at scale, and only big petrochemical has the infrastructure at scale to be part of the solution.

Outcomes – Carbon capture (hundreds)

Combining search terms such as MOF and COF (metal/covalent organic frameworks), carbon capture, storage yields a total of a few 100 papers, with chances of false positives making the counts imprecise. Carbon capture is clearly not a big theme for Green Chemistry. This might be a healthy recognition that many carbon capture schemes stand zero chance of making an impact at scale (the laws of thermodynamics are against them) or that other journals are more appropriate for publishing groundbreaking work.

Outcomes – Catalysis (4099/6315)

Who can argue against the need for better, greener catalysts? The numbers show that the Green Chemistry community believes that the planet will be saved by having better catalysts.

I have no way to analyse these catalysis papers. Maybe we now have many large-scale catalyst systems using greener chemistries, at lower temperatures. Or novel catalysts allow more carbon-efficient syntheses, a key green metric. One purpose of this paper is to stimulate debate. The data are available in the Excel sheet; those with expert knowledge in this area are encouraged to mine the data to help to appraise the impact of these many papers. I would be delighted if the impact on the planet turns out to be significant.

Outcomes – Life Cycle Analyses 90/230

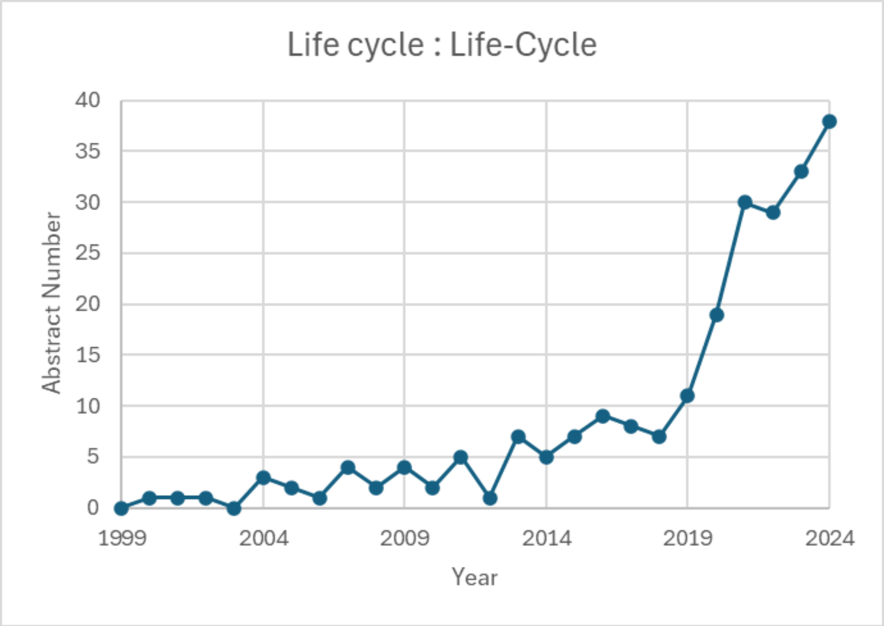

This graph of references to life cycle analyses (LCA) is plotted as absolute numbers, so the increase after 2018 from a low level is partly due to more papers being published. But it shows that the Green Chemistry community has been rather late to the idea that gauging the impact of a “green” product requires a holistic view of the upsides and downsides compared to alternatives.

But LCAs are complex and cumbersome and don’t make it easy to play intelligent “what ifs”. So since 2019, I have provided (free, open access) Open Analyzer (OA) (https://www.stevenabbott.co.uk/Open-Analyzer/Analyzer.php) which allows anyone to create an LCA-Lite, that others can explore in their browser, critique, clone, create alternatives and generally get a feel for the likelihood that a specific approach will be good, bad or neutral for the planet.

The screenshot shows an example from OA looking at a typical bio-based alternative to a petrochemical:

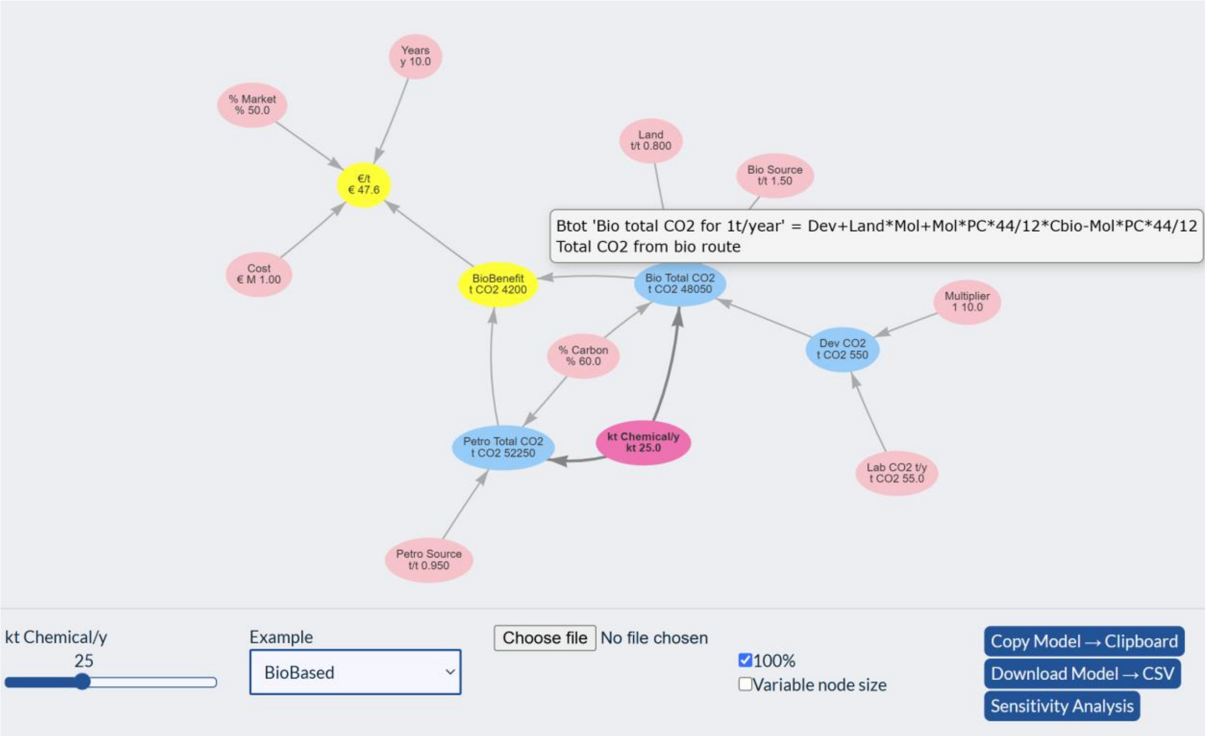

Btot 'Bio total CO2 for 1t/year' = Dev+Land*Mol+Mol*PC*44/12*Cbio-Mol*PC*44/12
Total CO2 from bio route

kt Chemical/y
25

Example
BioBased

Choose file No file chosen

☑100%
☐Variable node size

At the moment the amount of chemical produced per year is highlighted and the slider can change it from its current value of 25 kt/y. The impact of moving to biobased is shown in yellow and the effects of changing each of the inputs in pink can be quickly found, allowing "what ifs" to be explored. Moving the mouse over any blue or yellow oval (in this case calculation of the total CO_2 via the bio route) gives you the equation behind it. Anyone who disagrees with the input values can change them, and those who disagree with the model can Copy or Download the model and put it into their favourite spreadsheet to produce their own version – and then Choose File to load it.

This example, based on numbers validated by an LCA expert in early 2019, shows that if, after 10 years of hard (resource consuming) work the biobased alternative sells 25 ktpa, then the saving is 4.2 ktpa. That is equivalent to 3.4 seconds of current global CO_2 production (a handy calculator is provided in the Excel workbook). Which raises the question of whether those 10 years and all those resources might have been devoted to something in the Mtpa (14 minutes) or, rather, to what really matters to something in the range of Gtpa (10 days).

At the time of producing that model, getting hold of relevant numbers required the expert input of a leader in LCAs. In 2025 the situation is different as we now *can* get hold of large amounts of curated CO_2 equivalent values via the Idemat_2024 (or latest year) Excel dataset from https://www.ecocostsvalue.com/data-tools-books/. It is interesting to look at some entries in the Idemat dataset that compare bio-based with petrochemical polymers and chemicals. Usually the differences are small and sometimes the bio-based is worse.

Returning to ACR, it is certainly the case the LCAs show problems with the CO_2 emissions and waste disposal issues of ACR. But now do the LCAs on the alternatives. Any process that can scale is going to have downsides. If there were a greener alternative easily available, we'd all be using it.

If, as the evidence suggests, the Green Chemistry community was not seriously engaging with the life cycle implications of their work till 20 years had passed, that is clearly sub-optimal. And if, as is the case, an LCA identifies downsides for any at-scale process such as ACR, then that is an open challenge to the community to find ways to make the process better. Similarly, if a green proposal turns out to be either bad for, or insignificant to, the needs of the planet, maybe it is greener not to use precious resources working on that project.

There's more to being green than CO_2

Other green metrics can be invoked. Whether Green Chemistry can celebrate major wins through reduction in water use (0/8), eutrophication (2/17), toxic waste (0/9), harms to public health (0/9) or other metrics is beyond my expertise, though the numbers of relevant papers, from those simple search terms, seem small. Again, the dataset exists so those with the appropriate knowledge can assess the outcomes.

Alternatives that require less complex safety infrastructure, less waste disposal costs along with higher throughputs would be welcomed; not so much because they are green but because they are just better all-round. The U.S. EPA Green Chemistry Challenge (https://www.epa.gov/greenchemistry/information-about-green-chemistry-challenge) shows progress across many years on 3 metrics:

- 830 million pounds of hazardous chemicals and solvents eliminated each year—enough to fill almost 3,800 railroad tank cars1 or a train nearly 47 miles long;
- 21 billion gallons of water saved each year—the amount used by 980,000 people annually;
- 7.8 billion pounds of carbon dioxide equivalents released to air eliminated each year—equal to taking 770,000 automobiles off the road.

The awards celebrate the work of large companies along with work from academics and small businesses which will tend to have smaller impacts. Although "small is beautiful" is a pleasant green-sounding mantra, if we want to save the planet it needs a megacorp to make big stuff happen.

However, even though the EPA work involved megacorps, the 7.8 billion pounds of carbon dioxide each year needs to be put into context. That's 3.5 Mt, which is nice, but is 50 min of annual human CO_2 emissions. It might be a good green practice to quote minutes of CO_2 equivalents; 7.8 billion pounds sounds much more of an achievement than 50 min.

"Every little helps"

This comforting phrase excuses using resources on kt CO_2 issues that are irrelevant to the world's current crises. But it is a false excuse. "Every little" takes precious financial and intellectual resources away from things that matter at scale.

"But you can't create scientific breakthroughs to order"

That is true. Curiosity-driven science can be its own good. So shouldn't we just fund green chemistry in the hope that a bit of curiosity leads to a breakthrough no one could have foreseen? Of course we need *some* brilliant minds exploring new horizons, but identifying and supporting such science is a difficult challenge.

It is hard to search through 11,000 papers for examples of potentially game-changing breakthroughs. It is entirely possible that I have missed them. It is equally hard to search for heroic failures – where there's a negative outcome to an imaginative tough challenge and where the lack of success is because science is hard.

If others can find examples of breakthroughs or heroic failures, these should be celebrated. Again, the dataset is there for others to mine.

Citation Mapping

The amazing VOSviewer (https://www.vosviewer.com/) made it possible to create a citations cluster map for Green Chemistry:

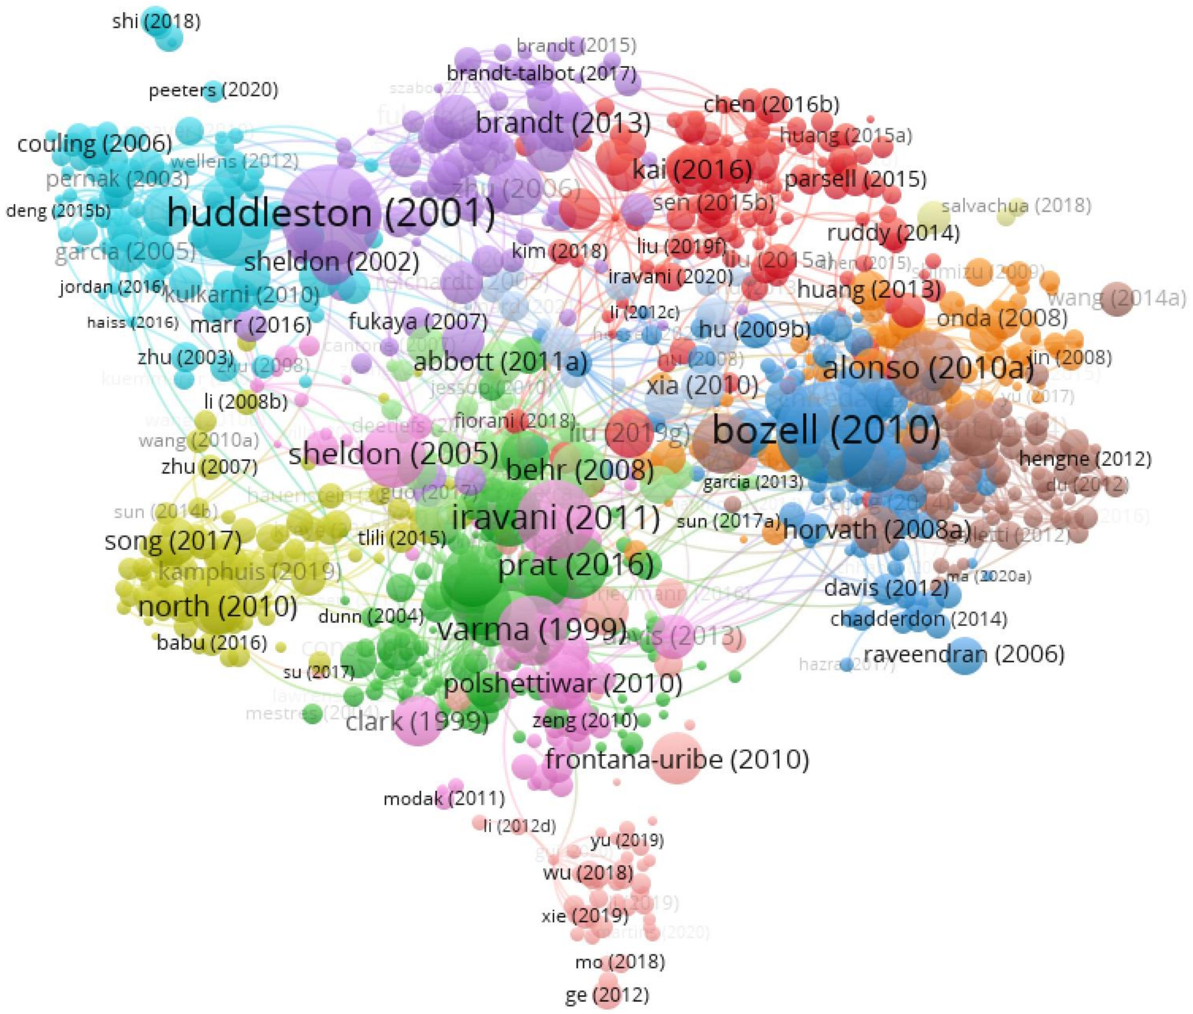

Although in principle I could do something interesting with this mapping approach, it is noted here as another possible way to look at the 25 years of Green Chemistry. For those with VOSviewer I am happy to provide the data files that can be fed into it. VOSviewer allows native download from WebOfScience but it takes 22 downloads to accomplish.

The future for green chemistry

Chemistry has an obvious key role in helping to save the planet. But in order to be as impactful as possible, the green chemistry community needs to be more self-critical. Even if the community avoids "greenwashing", in addition to consideration of the likelihood of success of any given project, the big question is whether, if successful, it could make a tangible difference at scale. If not, shouldn't these precious intellectual resources be devoted to more significant targets?

In terms of the key green chemistry target of CO_2 reduction, any project that is likely to work, at best, at the kT scale is not going to help. We need multi-Mt and preferably Gt solutions.

Applying this filter (along with filters on other outcomes the community deems to be important) would have changed the ethos of Green Chemistry. Even a cursory attention to the scale of the likely outcomes would have allowed authors, funding agencies and reviewers to note the contradiction between the high principles of green chemistry and the modesty of the likely outcomes.

So where are the global challenges that chemistry can help solve?

Ironically, agriculture is one of the major sources of CO_2. Rather than using agriculturally derived products in a bid to solve the world's problems, chemistry can identify ways for agriculture to be less harmful to the planet. Obviously greener ways to ammonia (80/170), presumably from green hydrogen(3/19), are one such focus. Chemistry's role in balancing agricultural productivity with long-term environmental health, could benefit from honest discussions on what chemistries would make a difference at scale.

Cement (54/448) is a major producer of CO_2, and the counts show some significant interest. Although I've not checked these papers specifically, I have some familiarity with the cement industry. It is dismaying to note that there is plenty of research saying "Cement creates too much CO_2 so here is work which uses an alternative raw material which would reduce emissions by X% per ton of cement". Unfortunately, it is obvious that their proposed waste/cheap raw material is available only in quantities that are irrelevant to the challenge. Lack of focus on things that can make a difference at scale is sadly common across disciplines.

When it comes to chemicals (not to be confused with "the oil industry"), these are relatively small contributors to global warming, so it is important to make an intelligent choice of targets for green chemistry.

Heretical as it might sound, efficient use of crude oil and its ACR equivalent, might well be optimal use of resources, especially as crude oil is not only "natural" but doesn't eat into food resources.

For this, Advanced Chemical Recycling (in its most general form) needs further work. Of course it is currently not very successful. Of course it has many challenges. But the (often well-intentioned) resistance to ACR is best overcome via smart chemistry that gives more efficient conversion back to virgin feedstock, with increasing ability to handle complex recycled mixtures and the waste necessarily generated in the early days.

There are some smaller specific wins (at modest Mt scale) which are nice to have. Arguably PET is well on the way to being efficiently recycled via its own unique chemistry, with a key problem (in many countries) being access to well-sorted PET bottle feedstock and then to equivalent quality of textile fibres, if the problems of what to do with the co-cotton can be solved.

Chemistry cannot solve irrational national or regional recycling strategies, but it can ensure that for any given large-scale recycled input, the best chemistry is available to handle what is being provided.

What about green chemistry's role in reducing harmful chemicals? One approach is to focus on reducing exposure to the harms of *any* chemical. In an industrial context that is already "mission accomplished". Health and safety regulations have meant that harms from "toxic chemicals" are negligible inside well-run manufacturing facilities. The counter examples, such as PFAS, show that well-run safety checks inside large corporations flagged up the dangers relatively early. It was management decisions that overruled the scientists, and brave whistleblowers who finally exposed the issues.

But chemical facilities produce "toxic waste". Here we have two choices. We can wait years to replace a toxic chemical with a drop-in replacement that can be flushed straight into the ecosystem. The evidence from 25 years of Green Chemistry is that this is not a successful strategy. Or we can provide cost-effective systems (in the broadest sense) that deal with toxic waste. But there are only so many chemistry resources available. Let's agree on key toxic waste problems and use lean, green chemistry to solve them, at scale, sooner rather than later.

It is likely that some of these solutions will involve enzymes or microbes, so partnership with our more biological colleagues will be vital. Adopting a "benign by design" approach is good if the team working on it first thinks through whether the effort (and timescale) involved is commensurate with the harms it will reduce compared with the upsides of alternative approaches.

To be green is to save the planet sooner rather than later, and this can only be done by being honest and realistic about the relatively few issues, at scale, that need to be tackled.

My in-depth encounter with 25 years of Green Chemistry (and I urge readers to access the dataset for themselves, to draw their own conclusions) has confirmed what I have suspected for many years. The community has misapplied precious resources on issues of small or negative impact rather than attempt to be truly green and solve problems that matter at scale.

Let's make the next 25 years of Green Chemistry, and green chemistry in general, ones where good science is devoted to solving easily-identified difficult problems, at scale, making the planet significantly better than it would have been without this work. To say it one final time, good intentions are not the same as good outcomes, and rather than waste precious research resources on insignificant outcomes, let's focus resources on things that matter.

Acknowledgements.

I normally acknowledge all those who help me with a project. But although I have received considerable help from some very knowledgeable people in the green chemistry community, it is inappropriate to risk their names being associated with what many might see as an attack on the green chemistry world. I have, instead, thanked them personally.