UvA-DARE (Digital Academic Repository)

Catalysis, God's Algorithm, and the green demon

Rothenberg, G.

Publication date

2009

Document Version

Final published version

License

CC BY-NC-ND

[Link to publication](#)

Citation for published version (APA):

Rothenberg, G. (2009). *Catalysis, God's Algorithm, and the green demon*. (Oratiereeks; No. 337). Vossiuspers UvA.

General rights

It is not permitted to download or to forward/distribute the text or part of it without the consent of the author(s) and/or copyright holder(s), other than for strictly personal, individual use, unless the work is under an open content license (like Creative Commons).

Disclaimer/Complaints regulations

If you believe that digital publication of certain material infringes any of your rights or (privacy) interests, please let the Library know, stating your reasons. In case of a legitimate complaint, the Library will make the material inaccessible and/or remove it from the website. Please Ask the Library: <https://uba.uva.nl/en/contact>, or a letter to: Library of the University of Amsterdam, Secretariat, Singel 425, 1012 WP Amsterdam, The Netherlands. You will be contacted as soon as possible.

UNIVERSITY OF AMSTERDAM

Gadi Rothenberg

Catalysis, God's
Algorithm, and
the Green Demon

Faculty of Science

Catalysis, God's Algorithm, and the Green Demon

Vossiuspers UvA is an imprint of Amsterdam University Press.
This edition is established under the auspices of the University of Amsterdam.

This is inaugural lecture 337, published in this series of the University of Amsterdam

Cover design: Nauta & Haagen, Oss
Lay-out: JAPES, Amsterdam
Cover illustration: Carmen Freudenthal, Amsterdam

ISBN 978 90 5629 589 9
e-ISBN 978 90 4851 112 9

© Gadi Rothenberg, 2009

All rights reserved. Without limiting the rights under copyright reserved above, no part of this book may be reproduced, stored in or introduced into a retrieval system, or transmitted, in any form or by any means (electronic, mechanical, photocopying, recording, or otherwise), without the written permission of both the copyright owner and the author of this book.

Catalysis, God's Algorithm, and the Green Demon

Inaugural lecture

delivered on the appointment to the
chair of Heterogeneous Catalysis and Sustainable Chemistry
at the Van 't Hoff Institute for Molecular Sciences,
University of Amsterdam
on Friday 29 May 2009

by

Gadi Rothenberg

VOSSIUSPERS UvA

Mevrouw de Rector Magnificus,
Mijnheer de Decaan,
Dear friends and colleagues,

The Chinese have a curse, that reflects the differences between eastern and western cultures. It does not invoke the cursed person's bodily functions, and it makes no assumptions as to the occupation of their mothers or grandmothers. It simply says: 'May you live in interesting times.' But we, here in the west, we *like* to live in interesting times. Makes life interesting. And happily, the coming years look very interesting for science and for scientists. We're making one of our periodic historical comebacks, riding on the wave of the technological advances of the last three decades, and getting an extra boost from hot issues such as global warming and the energy crisis. Of course, with the current financial crisis you would also expect religion to make a comeback, filling the churches every Sunday morning. I mean, just look at how many people are here in this beautiful Lutheran church, and it's not even Sunday, it's Friday afternoon – these Lutherans must be doing something right.

Uninteresting times also have their advantages. When one year is much like another, people can afford to take the long view and invest in infrastructure and other long-term projects. The mere fact of knowing that something will not change is an important asset. For example, one of the most attractive aspects of solar and wind energy is that you know both the current price of the raw material (zero) and the future price of the raw material (also zero).

If we look at physical chemistry, then thermodynamics is all about taking the long view. For thermodynamics, time is not important. A piece of diamond will eventually turn into piece of graphite. As a former soldier, I can tell you that one of the things that makes life bearable is the knowledge that eventually your three years will pass, and you'll be a civilian again. Alternatively, if you stay long enough, you'll be promoted to sergeant-major. In fact, since sergeant-majors were not immensely popular among us common soldiers, we used to say that if

you would tie a donkey close to the mess room, it would be promoted to sergeant-major within five years. In fact, we could even speculate here that any student, who takes the long view and stays long enough at university, will eventually be promoted to full professor, but I guess more rigorous studies are needed to prove this hypothesis.

Young people are the ones that can actually afford to take the long view. Many of them do this by ‘investing in their future’, taking for granted that they will in fact have a future. When you get older, you realise that you cannot take the future for granted. You may have a future ahead of you, but on your way to this future you must make the most of the present. Curiously, the advice that older people often give to youngsters is to ‘take the long view and invest in your future’, even though they know that enjoying the present is just as important.

All that being said, young people are often impatient. Things may work out if you wait long enough, but you are not happy with this waiting. Instead of waiting for things to happen, you try and make them happen faster. And if possible, you try to change things without changing yourself in the process. That is, things will happen, things will improve, but **you** will stay the same, and therefore **this you** will enjoy the new, improved situation. Strikingly, this is one definition of a catalyst: something that makes things happen faster, while remaining itself unchanged at the end of the process. So in many ways, an impatient young man is a catalyst.

For me, this obsession with time portrays the main difference between catalysis and thermodynamics. It’s all about time, about making things go faster. More specifically, for us scientists this mirrors the concern about the time we have for solving a given problem. Solving complex problems is vital for scientists. It’s how we justify our existence as academics. Because being an academic is quite a special situation. You’re usually surrounded by fellow academics at work, so much that you’re in danger of forgetting the world at large. But every so often, you do get a chance to meet people outside the ivory tower, and when I do, the conversation sooner or later gets to the ‘so what do you do?’ stage, and I say (and this is invariably in Dutch)

- ‘Betaal je je belasting wel?’
- ‘Ja!’
- ‘En ik geef het uit.’

This, in fact, is what we do. We spend the taxpayers’ money. We may complain that we don’t get to spend as much of it as we used to, but the bottom line is that

CATALYSIS, GOD'S ALGORITHM, AND THE GREEN DEMON

we do spend it, and in one way or another, we are held accountable for it. We are supposed to do something useful, in the broadest sense of the word, for society. Hopefully, we are solving some important problems, and God knows that we're not short of important problems to solve.

When the questions you want to solve are complex and require years of investment, the time in question is really the time of your lifetime. **How much time do you still have?** If you look at your left hand, you will see a curved line that starts about 2 cm under the base of your index finger and curves clockwise towards your wrist. This, according to chiromancy, is your life-line, and the longer it is, the better. And this is also all that I know about chiromancy, but still I like it a lot, because the line on **my** left palm is a continuous curve that goes all the way to the back of my hand, and the one and only time that someone who professes to know about these things looked at it, they were astounded and said that according to this, I will live forever. So I decided to quit while I'm ahead, and not go further into this subject.

But going back to our catalyst analogy, we see that in reality things usually change step-by-step, just like in a chemical reaction, and as a catalysis chemist, you study these changes, and you seek **that special catalyst** that will give you the optimal pathway. The catalyst that will give just the product that you want (what we chemists call high selectivity), in quantitative yield (so a kilo of product for every kilogram of reactant), and do it fast – because time is of essence. And not only in the research lab. The biggest headache in the chemical industry is related to reactor space-time yield, meaning that if you can run your reaction faster, and with high selectivity, you can use a smaller, cheaper, and safer reactor. For large-scale processes, you're talking millions and millions of euros in reduced costs, not to mention the benefits to the environment.

One of my favourite examples is the manufacturing of aniline. This is a rather simple molecule – an ammonia molecule where one of the hydrogens is swapped for a benzene ring. Yet it is the namesake and key original product of the largest chemical company in the world: the Baden Aniline and Soda Factory, or in German, **Badische Anilin- und Soda-Fabrik**, more commonly known by its initials, BASF. Aniline is used for making dyes, polymers, pharmaceuticals, you name it. It's the simplest aromatic amine, and in principle you should be able to make it by reacting benzene and ammonia, getting aniline and hydrogen. This is a simple reaction, at least on paper. Moreover, it is thermodynamically allowed –

there is no thermodynamic problem here. All you need is to find the right catalyst. But this much is easier said than done. In practice, the manufacturing of aniline is an ugly and wasteful three-step process, where you first take the ammonia, oxidise it to nitric acid, and then take this nitric acid and mix it with benzene in concentrated sulphuric acid to give nitrobenzene, and then use hydrogen to reduce this nitrobenzene back to aniline. So your starting material is an amine, your product is an amine, and you go from one to the other in, to put it mildly, a rather round-about route.

Two years ago, I visited Ludwigshafen as a consultant to BASF's catalyst discovery program. We started talking about aniline, and I asked the section heads at the ammonia lab (just to give you an idea, BASF has three research labs in Ludwigshafen, each one with about 2000 researchers; roughly 500 PhDs and 1500 technicians; for comparison, the entire Faculty of Science at the UvA employs fewer than 400 PhDs and 200 technicians); anyway, I asked the section heads at the ammonia lab what's the situation with the direct catalytic route from benzene and ammonia to aniline. They said that every couple of decades, someone thinks that they're going to find the optimal catalyst for this reaction, but it hasn't happened yet. But, last October one of these guys said to me at a conference that he thinks that this time they found it. So who knows, maybe by the time I give my retirement lecture, hopefully sometime in the spring of 2034, this example may be obsolete. Set the date in your agendas.

By and large, this is the type of research we're interested in. The people in my group are searching for new catalysts that can improve existing chemical processes and hopefully enable new reactions. Finding such catalysts is difficult. In fact, it is so difficult, and lucky breaks happen so seldom, that you start asking yourself, are we doing things in the right way? Is there no better method for finding new catalysts? Maybe we could find good catalysts intuitively – just come up with the right ones without getting hopelessly lost in all the dead ends?

Imagine if we could do that. Imagine if we had an algorithm that would instantly give us the formula of the optimal catalyst for a given reaction. In fact, imagine how things would be if we had an algorithm that would instantly provide us with the optimal solution to any given problem. This, ladies and gentlemen, is God's Algorithm. You can think of it as a table of questions and answers, one answer per question, no more. And we mortals, of course, have no access to it, although once we almost had it, as I will show you now.

CATALYSIS, GOD'S ALGORITHM, AND THE GREEN DEMON

To understand the background to God's Algorithm we must turn to the Bible, and more precisely to the first three pages of the Book of Genesis. It all starts in the Garden of Eden, which was a very lush and tropical place. In Chapter 2, verse 9, we learn that God grew there all the trees that were nice to look at and had good fruit to eat, and two trees are especially mentioned: the tree of life and the tree of knowledge. And in Chapter 2, verses 16 and 17, God gives man a simple and unambiguous directive:

טז וַיֹּצֵא יְהוָה אֱלֹהִים עֵץ-הָעֵדֶם לְאֹמֶר מִכָּל עֵץ-הַגֶּן אָכַל תֹּאכְלִי;

יז וּמִעֵץ הַדַּעַת טוֹב נָרָע לֹא תֹאכַל מִמֶּנּוּ כִּי בְיוֹם אֲכָלְךָ מִמֶּנּוּ מוֹת תָּמוּת

[Genesis 2:16-17]

So man was forbidden the fruit of only ONE tree: the tree of knowledge. This is an important point – God did NOT forbid man to eat from the tree of life, only from the tree of knowledge. Afterwards, God created woman, and there was the story with the snake, and both man and woman ate from the forbidden fruit of the tree of knowledge, and God got quite upset with them. But the really interesting thing is that this was not the direct reason that Adam and Eve (who by the way were named Adam and Eve only AFTER they ate from the tree of knowledge) were expelled from the Garden of Eden. No. The real reason is given by God himself one page later, in Chapter 3, verse 22:

כב וַיֹּאמֶר יְהוָה אֱלֹהִים; הֵן הָאָדָם הָיָה כְּאֶחָד מִמֶּנּוּ לְדַעַת טוֹב וָרָע וְעַתָּה; פֶּן-יִשְׁלַח יָדוֹ וְלָקַח גַּם מִעֵץ הַחַיִּים וְאָכַל וַחַי לְעֹלָם;

[Genesis 3:22]

And this literally confirms what the snake promised to Eve, that when she and Adam would eat from the tree of knowledge, they would be equal to God in their ability to distinguish between 'good' and 'bad', which is the ultimate classification analysis ability. And now God is concerned that if man should reach out and take the fruit of the tree of life and eat it, he shall live forever. And so God takes drastic actions, casts Adam and Eve out of the Garden of Eden, and places the angels and flaming sword to see to it that they don't get back in.

Now, let's get back to God's Algorithm and see if we can understand why God was so concerned about Adam and Eve possibly eating also from the tree of life. Let's assume that God's Algorithm is indeed some kind of matrix, wherein each question has one optimal answer, and where the access to the answer is instantaneous. So when you see the question – you immediately KNOW the answer. Think for a moment – did this ever happen to you? It happens to all of us occasionally when we try and solve problems. Most of the time, we have to think about the answer, but sometimes we KNOW the answer immediately, it 'comes in a flash', intuitively. Is this intuition a manifestation of God's Algorithm? How does intuition work? I'm not sure, but there is some connection between intuition and experience. When you encounter a situation that is similar to a 'previously solved' situation, you can make an 'educated guess' at the answer. Many of you have kids, and remember their first weeks of walking. A little child learning to walk will stumble often, unsure of where to place his legs, especially if the terrain is unfamiliar. But as adults, we hardly ever stumble. In fact, we hardly ever even think about the process of walking, even if this is a new street where we've never walked before. Our mind simply compares this situation to a host of common factors from countless previous similar situations and immediately produces a solution that is a very good 'educated guess' on how to walk on this new pavement. This applies not only to walking, but to every type of problem that we frequently encounter.

As we get older, we gain more and more experience, and so we can solve more and more problems intuitively. But there's a catch, of course, because as we get older, our bodies and our brains deteriorate, and we tend to forget things. We do not have total recall.

Now imagine how things would be if man had also eaten from the tree of life. He would then live forever. His brain would not deteriorate, and it's a fair bet that he would have total recall. And since according to God himself, the act of eating the fruit of the tree of knowledge has already given man the god-like ability to distinguish between 'good' and 'bad', the fruit of the tree of life would have given mankind full access to God's Algorithm. Small wonder that God was concerned.

Fortunately, God cast Adam and Eve out of the Garden of Eden BEFORE they sampled the fruit of the tree of life. I say 'fortunately' because as a scientist, it would have been extremely boring if everyone knew immediately the optimal answer to every problem. But you must admit that a problem-solving approach

CATALYSIS, GOD’S ALGORITHM, AND THE GREEN DEMON

founded on knowledge-based intuition is attractive. Moreover, it works – people apply this approach daily, some more consciously than others. One small and rather special group of people studied in this connection is professional chess and go players. Both of these board games are based on a small set of simple rules and always start from a single, well-defined position. Yet each game quickly branches out, offering a mind-boggling number of possible tactics and moves. Interestingly, research on high-ranking chess players showed that they don’t ‘see’ more moves than lesser players. But the moves that they see are ‘good’ ones. Somehow, their brain simply ‘filters out’ the ‘bad’ moves.

So the question is, can we apply this intuition-based method in some way to solving problems in catalysis? More specifically in my case, could I formulate a working algorithm that would increase the chance of discovering new catalysts and optimising existing ones? And will it actually work in practice? These questions formed the basis of my application to the position of Assistant Professor in the Department of Chemical Engineering here at the UvA, eight years ago. Today, I can proudly tell you, even if I do say so myself, that the answers are yes, yes, and yes.

So how do we do this? How do we predict what I call the Figures of Merit for new and existing catalysts? The approach is analogous to the prediction of the activity and efficiency of pharmaceuticals, using quantitative structure-activity and structure-property relationship models. But, once you start with this analogy, you realise very quickly that there are two crucial differences:

The first is that catalysts are not drugs. They do not fit nicely into the category of ‘small-to-medium-sized molecules with one or two polar and/or hydrophobic groups’. In fact, in heterogeneous catalysis, the catalyst is most often a discontinuous, sometimes totally amorphous, solid.

The second aspect that differs catalyst optimisation from drug development is the range of reaction and process conditions. You will all agree that humans are much more complex than inorganic materials, but in vivo processes are rather constrained as far as temperatures and pressures go. All the reactions in our bodies run at roughly 37°C and 1 atmosphere pressure, in an aqueous solvent (since we’re about 90% water). In contrast, industrial catalytic processes operate over a wide range of temperatures and pressures, from a chilling –40°C in the production of many pharmaceutical intermediates, all the way to +800°C in fluid catalytic cracking reactors. These extreme process conditions often cause dramatic changes

in the catalyst, meaning that the actual material that is catalysing your reaction is very different from the stuff that you packed into the reactor bed three weeks ago. And all this basically leaves you two options if you're interested in structure/activity relationships. Either you follow the exact catalytic cycle under so-called operando conditions. This means finding a suitable method for following the reaction, and usually building very sophisticated equipment for studying catalytic reactions exactly as they happen. Or, you try and make some very good guesses as to the connections between the material you started off with and the actual catalyst. Or maybe a little of both...

What we actually do in my group is implement a human-sized version of God's Algorithm to find and optimise catalysts. When I say human-sized, I mean that the algorithm will only solve **specific** problems, and that our expectations of it must be realistic. In fact, this setting of realistic expectations is an extremely important point that pertains to all research projects, and especially to those that involve computer models. Setting realistic expectations means that in **some** cases, providing that you start from good experimental data, and you have a good understanding of the reaction system, and sufficient validation experiments, the model will show you the way to the optimal catalyst and reaction conditions. In most cases, however, the best that you can expect is that the model will point you away from 'bad experiments'. **But make no mistake – this is already a huge improvement!** Imagine how nice it would be if you could avoid doing 'bad experiments' – just like Gary Kasparov avoids 'seeing' the bad moves in a chess game.

We do have one extremely important advantage today: cheap and powerful computers. Using computers, our models will have total recall, because computers (unlike the scientists using them) rarely forget anything. Thus, if we can translate our catalysis problem into something that a computer can understand, we're halfway there. To do this, however, we must first understand the problem of catalyst optimisation, and especially define the space in which we want to optimise our data. To connect catalyst structure with catalyst performance, we redefine our system using three multidimensional spaces, that I like to call A, B, and C. Space A is a grid containing all the catalyst structures; space B, the descriptor space, contains the values of the catalysts' descriptors, and space C contains the values of the figures of merit. The descriptors can be internal parameters, such as backbone flexibility, partial charge on the metal atom, polarity, surface area, or crystallite size. They can also be external parameters such as temperature, pres-

CATALYSIS, GOD’S ALGORITHM, AND THE GREEN DEMON

sure, or solvent type. All of these parameters may influence the reaction outcome, and a model may include both internal and external descriptors. The figures of merit in space C are what we’re after, e.g. the catalyst turnover number, the turnover frequency, product selectivity, price, and so forth.

By re-defining the system this way, we translate our abstract problem in catalysis into a problem of relating one multi-dimensional space to another. It’s still an abstract problem, but now we can use our experimental data to build a link between the figures of merit in space C, and the catalysts for which this data was measured, in space A. We can go from space A to space C in two steps: first, going from A to B, by computing molecular descriptors. Then, going from B to C, via quantitative structure-activity/structure-property relationship models.

So what we do first is collect experimental data, then build a descriptor model, then model part of the experimental data, and use the other part for validating the model. We can then predict the figures of merit for new catalysts, which we have not yet tested, and then synthesise some of those and validate our predictions.

This is one iteration. We then rebuild the model based on all the data that we now have, and make predictions for a new set of catalysts, and then synthesise some of **those** and validate our predictions, and repeat the whole process. With each iteration, we improve the model, until we reach a catalyst that we’re happy with.

One key advantage of this iterative approach is that you can use the computer to ‘mix and match’ different catalysts and reaction conditions. In this way, you can effectively search in a space that is about 20 times larger than your experimental space. For example, if your time and budget allow for the synthesis and testing of only 50 catalysts and reaction conditions, you can use our models to search in a space of 1000 catalysts and reaction conditions.

The key to translating the catalysis problem into something that a computer can understand lies in space B, the descriptors space. Find good descriptors, and your problem is over, because the statistical methods themselves are well known. Doing this for homogeneous catalysts, which are molecules or organometallic complexes, is rather straightforward, because the catalyst is usually well defined, so I started there. Actually, I simply applied the teaching of David Avnir, who explained to me in my first year at university that one should ‘always start with the easy part’. So let’s see how we build descriptors for molecules and complexes.

Catalyst descriptors can be calculated at several levels. Three-dimensional descriptors, which are based on optimised geometries, can be calculated using molecular mechanics force fields and quantum mechanics calculations. These 3D descriptors offer a realistic representation of chemical systems. The problem is that their computational cost depends on the system's size. If you plan to optimise large numbers of catalysts, 3D descriptors are simply too expensive. In such cases, we turn to the simpler two-dimensional descriptors (also called topological descriptors). These are derived directly from molecular connectivity tables, without using any 3D atom coordinates. Topological descriptors are typically **five orders of magnitude faster** than 3D descriptors (that's 100,000 times faster). But as in many other things in life, you get what you pay for. The lower cost is offset by several limitations: First, topological descriptors have no direct heuristic interpretation, because they are far away from our 'chemical intuition'. Second, 2D descriptors neglect conformational information, and because they are two-dimensional, they cannot be used for modelling chirality. Still, if you're planning to model large numbers of catalysts, topological descriptors are your only option. The models are then 'rough and fast' or, as I prefer to call them, 'quick and dirty'. This type of modelling is very appealing to me, since it's very similar to 'shake and look'-type experiments, which is the only type of experiments that I'm good at.

With the help of some highly talented chemometricians and mathematicians, namely Hans Boelens, Jos Hageman, David Iron and Johan Westerhuis, and one gifted PhD student, Enrico Burello, we put together a working piece of software that could accelerate the optimisation of homogeneous catalysts. We even proved mathematically that this approach would increase the chances of finding optimal catalysts. All we had to do was demonstrate it in a real-life environment. But, you must remember that I was a mere postdoc, so nobody wanted to apply this. Companies said 'not interesting', or 'show it to me when you've found some catalysts', or even, quite conceitedly, 'oh, yes, we already know how to do this' (but I'm not going to tell you which company said that). But, happily, there was one key player in the chemical industry who said 'let's try this'. This was the French company Rhodia. They were the first to invest in our predictive modelling approach for their nickel-catalysed hydrocyanation process for making nylon precursors. Today, this collaboration has resulted in several papers and patents on new catalysts, and Ana Maldonado from our group will join Rhodia next month, to spread the message of predictive modelling in the Rhodia Lab of the Future in Bordeaux. Of

CATALYSIS, GOD'S ALGORITHM, AND THE GREEN DEMON

course, now that we have demonstrated that the concept works, several other companies are funding projects in the group, but it all started with Rhodia, and I am very happy that Patrick Maestro, head of the Lab of the Future, is here with us today. Chapeaux, Patrick.

So we're in a good position to apply our small-time version of God's Algorithm to homogeneous catalysis. But, as you may have noticed on the invitations, I am supposed to be the Chair of **Heterogeneous Catalysis and Sustainable Chemistry**, and while the loose boundaries of sustainable chemistry probably allow for some forays into homogeneous catalysis, it would be nice to try and apply the same methods also in heterogeneous catalysis.

The problem is that finding good descriptors for heterogeneous catalysts is much more difficult. Unlike molecular catalysts and organometallic complexes, the activity of solid catalysts depends on a multitude of parameters: different types of active sites, synthesis conditions, thermal treatments, and ageing. Moreover, the properties of many solids can change discontinuously. New phases can form at different compositions, temperatures and pressures, and even the catalyst particle size can influence the reaction. Nano-sized gold particles, for example, are very different catalysts from 'bulk gold', and two supported gold catalysts can have different activities even when they contain identical amounts of gold and support, respectively. Solid surfaces are anything but uniform, and solid catalysts have a variety of sites. To complicate things further, sometimes the real active sites are not those observed in characterisation studies, but rather metastable defects that are difficult to characterise.

Because of this, relying solely on catalyst composition parameters is impractical (compositional descriptors are mainly applicable in cases where the catalyst is a crystalline material, and where the changes in composition do not lead to phase changes). Instead, we need a descriptor toolbox that can account for the discontinuities and nonlinear dependencies. This is a really big challenge.

So how are we approaching this problem? First, we're looking at a variety of reactions that ARE catalysed by crystalline materials, mostly ceramic mixed oxides that are also stable at high temperatures. The big advantage here is that we can directly connect the elemental composition to the final catalyst structure. This means that we can screen large numbers of possible catalyst compositions on the computer, and build models that will predict the activity of these virtual catalyst libraries. We can even use evolutionary algorithms to 'breed' new mixed oxide

catalysts with improved properties, just like our colleagues in the genetics department can design apples with a longer shelf-life, or cows that give more milk. The applications of such crystalline catalysts range from propylene manufacturing all the way to turning lignin, which is the gooey stuff that holds trees together, into valuable industrial chemicals. So this side of predictive modelling in heterogeneous catalysis is enough to keep us busy for the next few years.

On the more fundamental side, we're working with non-crystalline materials and supported catalysts. Here we're trying to build a simple model that will connect the catalyst preparation recipe with the final surface composition. It's still a far cry from predicting catalytic activity *in silico*, but it's an important start. For example, in a typical preparation recipe, you may take an alumina support, pour onto that a 50:50 solution of gold and silver, burn off any residues, and then reduce the surface with hydrogen. Will you get 50:50 gold:silver on the surface? The answer is no, because the segregation energies of gold and silver are different, and the cohesive forces between two gold atoms or two silver atoms differ from the adhesive forces between a gold atom and a silver atom. But if you know these segregation energies, you can model the distribution of the surface composition. And since heterogeneous catalysis occurs on the surface, you can also start predicting catalytic activity. So despite the fact that this is a serious challenge, I'm pretty confident that by 2034, and possibly even earlier, we will have some good working models.

So whichever way you look at it, the marriage between experimental catalysis and predictive modelling is a beneficial one. On one hand, we are developing new theories and models, that open opportunities and new pathways for basic research. Simultaneously, we are generating high-quality results, from new catalysts to automated workflows that are applied in the chemical industry. And as we all know, scientific research is all about getting results. Or is it? I mean, anybody can do experiments and get results. Maybe scientific research is really about the discussion, the ideas that stand behind the results?

Most of you would probably say that both are important, since scientific papers invariably contain both results and discussion. But since we all know that it's 'publish or perish' out there, it's crucial to understand how to compose our papers. I'm interested, are results more important for us in our publications? Or discussion? Is one more important than the other? We have a great many scientists here, so let's have a small poll:

CATALYSIS, GOD'S ALGORITHM, AND THE GREEN DEMON

Raise your hand high, if you think that for us, as scientists, it's more important to publish results;

OK, now raise your hand high if you think that for us, as scientists, it's more important to publish discussion;

Then do the green demon experiment...

So we see that the coming years hold the promise of interesting times, with no shortage of fundamental and applied problems in catalysis for us to solve. And as long as it promises to be fun, we will keep on doing it, because, to quote the astrophysicist Eric Robert Schulmann, 'the purpose of science is to get paid for doing fun stuff'. And even if we're not getting paid so handsomely, we are certainly doing fun stuff.

But there are additional aspects to becoming a full professor and chair, apart from doing cutting-edge research. The most obvious of these is the fact that you're entitled to go to the university's gown-tailor, and have one of these beautiful wool and velvet gowns made to measure. This is a memorable experience, because the tailor's atelier is an extremely traditional family place of haute couture, where de heer Rhebergen has been making gowns for over twenty years, just like his father, and his father's father, and his father's father's father, and you're almost feeling like you've entered into a live Monty Python sketch. But the gown is amazing, and when you wear it you feel a little of the grandeur and tradition of 375 years of University of Amsterdam passing to you in some small way. There is even a rumour that the hat was by originally designed by Rembrandt van Rijn, and he certainly lived in the right period, at least.

A less-celebrated side-effect of becoming a full professor and chair is that you cannot hide anymore from all of the policy functions and discussions. You cannot say, 'I just work here, I'm a postdoc or assistant professor, let someone else chair that committee'. You're expected to do your bit, and help steer the University of Amsterdam towards the right path. Personally, I have not suffered from this too much, probably because I am a rather opinionated so-and-so, and thus people prefer keeping me out of such committees. I'm also reasonably happy, so I try and keep my mouth shut, like that Franciscan monk who kept his vow of silence until the cook forgot to add salt to the soup. But there is one trifling point that I will address here, and that pertains to the dwindling numbers of science students.

Europe needs more scientists. This is a fact. It is also the title of a surprisingly readable and lucid EU document, the High-Level Group report. The Report, pub-

lished in 2004, declares a sizeable shortage of more than 700,000 MSc and PhD scientists, proposing rather optimistically to fill this gap by December 2010. That's not going to happen. And since governments know that it's not going to happen, some people are taking drastic actions. And we should watch out! What worries me in particular is the universities' (and indeed the government's) wish that more people should complete their studies, and even more seriously, that a **higher percentage** of students must complete their studies successfully.

If Dutch students don't crowd the waiting lists for chemistry and physics, the problem will not be solved at the university level. It must be addressed far earlier, in Dutch primary and secondary schools. Happily, the Dutch government is doing something about it, by raising teachers' pay and providing new programmes that show kids how exciting science can be. This is a long-term solution. Lowering the entrance criteria or the exams level at university is NOT the short-term solution. Because Europe does NOT need more scientists. It needs more GOOD scientists. And if fewer young people apply to study chemistry and physics, it also means that fewer GOOD people apply to study chemistry and physics. When the university press office asked me for a short quote for this lecture, I said to them simply, '80% of all scientists are not in the top 20%'. I actually wanted to say that '80% of all **students** are not in the top 20%', but after thinking it over, I decided to change it to scientists, rather than risk offending any of our precious few students.

So what ARE we going to do in the short term? I will propose to you now a good and viable solution. This is not a joke. It is a real solution:

We must keep the standards of our university high and open our gates to any student from any nationality, from all countries in the world. Quality in the sciences should be our only selection criterion.

Of course, this would mean teaching all the courses in the faculty of science in English. This is not a problem. The English level of Dutch students is exceptionally high, and all of our faculty members can easily teach in English. Last year, I proposed this to the Dean of our Faculty, Professor Karel Gaemers, and we agreed to run a test case. I taught a five-credit first-year course in basic physical chemistry, everything in English. The students evaluated the course highly, and an overwhelming majority of the evaluation forms noted that the fact that the course was given in English was a plus.

But, you may ask, even supposing that we could attract two hundred excellent candidates per year for chemistry from China, India, or Russia, who is going to

CATALYSIS, GOD'S ALGORITHM, AND THE GREEN DEMON

pay for their tuition? Very simply, we should extend to these excellent students the same nice conditions that the Dutch government gives to the sons and daughters of Dutch taxpayers. We can set the selection criteria extremely high, but those that pass will not pay any tuition. The UvA will not be the first institute to do this, because this policy was instituted two years ago by MIT and Harvard, where excellent applicants are exempt from tuition fees.

Where will this lead us? Well, we will have a practically limitless supply of eager and exceptionally bright students. If we invest in these excellent foreign students enough and treat them nicely enough, they will feel welcome here, and stay in the Netherlands after their studies. They will make a substantial contribution to the Dutch and the European economy, and help us meet the goals of the High-Level Group Report. Of course, we must consider the possible consequences of our actions. Bringing such able and intelligent foreigners to The Netherlands, and giving them good conditions. Do we want this? Why, if we don't watch out, these foreign guys will eventually swamp our academic institutes, and become professors and chairs. Hmmm...

Coming to conclusion, I thank the University Board and the Faculty of Science for entrusting me with this chair. I also thank the colleagues who have helped our research group, especially Kees Elsevier and Rob Zsom, whose valuable trust and diplomatic advice helped reduce the brain drain from the Netherlands. Thank you, Kees and Rob. Finally, I am very fortunate to have two excellent technicians, Marjo Mittelmeijer and Paul Collignon. Much of our group's international success is owed to their expertise. Thank you, Marjo, and thank you, Paul.

Ik heb gezegd.