CWI Syllabi

CWI
P.O. Box 94079, 1090 GB Amsterdam, The Netherlands
Telephone + 31 - 20 592 9333
Telefax + 31 - 20 592 4199
WWW page http://www.cwi.nl/publications_bibl/

CWI is the nationally funded Dutch institute for research in Mathematics and Computer Science.

Proceedings of the forty-second European Study Group with Industry

Amsterdam, The Netherlands, 18 - 22 February 2002

G.M. Hek (editor)

CWI Syllabus 51

Study groups with industry are meetings where people from industry join forces with mathematicians to jointly tackle industrial problems. The first such meeting was held in the sixties at the University of Oxford. Nowadays study groups with industry are organised in many countries, e.g. Australia, the USA and several European countries.

Since 1998 study groups with industry are also organised in The Netherlands. In 1998 a study group was hosted by the University of Leiden, in 1999 by the University of Eindhoven, in 2000 by Twente University, and in 2002 the study group was hosted by the CWI and the University of Amsterdam. The next study group will take place in Leiden again, in February 2003. For up-to-date information please follow the links `www.wiskgenoot.nl/swi` or `www.stw.nl`. The 42nd European Study Group with Industry was organised by the University of Amsterdam and the CWI, the National Research Institute for Mathematics and Computer Science in the Netherlands. The programme was supported by the section *Industriële en Toegepaste Wiskunde* (ITW) of the *Wiskundig Genootschap*, with financial support from the technology programme *Wiskunde Toegepast* of the technology foundation STW and the EW (Exact Sciences) programme of The Netherlands Organisation for Scientific Research (NWO), and with additional support from the CWI, the University of Amsterdam, and the European Consortium for Mathematics in Industry (ECMI).

Organising committee:
G.M. Hek
M.F.M. Nuyens
M.A. Peletier
R. Planqué
H. van der Ploeg
G. M. Terra

2000 Mathematics Subject Classification:
00B25, 35QXX, 60JXX, 92D40, 92C05, 31AXX, 90B80, 90C10.

ISBN 90 6196 516 0
NUGI-code: 811

SWI 2002

pp. v – vi

Voorwoord

De positie van de wiskunde in Nederland heeft in de afgelopen maanden veel aandacht gekregen. Op 18 april 2002 presenteerden de wiskundige onderzoekscholen de nota *Nieuwe dimensies, ruimer bereik*, waarin gepleit wordt voor een versterking van het wiskundig onderzoek in Nederland, ondanks de teruglopende studentenaantallen. Die studentenaantallen vormen een bedreiging voor het voortbestaan van de negen wiskundeopleidingen die ons land rijk is. Dat de kwaliteit van de opleidingen over het algemeen goed is bleek uit het rapport van de onderwijsvisitatiecommissie voor de wiskunde dat op 25 juni 2002 werd aangeboden aan de voorzitter van de VSNU. Op dit moment voeren de universiteiten overleg over het Bachelor Convenant Natuurwetenschappen, waarbij de vraag aan de orde is of zelfstandige bacheloropleidingen in de wiskunde (en de natuurkunde en scheikunde) in stand zouden moeten blijven, of dat zij zouden moeten opgaan in bredere bachelorprogramma's.

Het NWO programma Wiskunde Toegepast heeft tot doel het toegepast wiskundig onderzoek te versterken. Dat gebeurt in de eerste plaats door de financiering van promotieonderzoek en post-doc posities. Maar het programma rekent het ook tot haar taak om activiteiten te ondersteunen die kunnen bijdragen aan een betere beeldvorming van het vakgebied. De Studiegroep Wiskunde met de Industrie is een jaarlijks terugkerend evenement dat een uitstekende bijdrage levert aan die beeldvorming en dan ook van harte wordt ondersteund door Wiskunde Toegepast. Elk jaar zijn het vooral jonge wiskundigen van de verschillende universiteiten, die het initiatief nemen voor de organisatie van de studiegroep. Daarmee laten zij zien dat de wiskunde een springlevend vak is, dat onmiddellijk toepasbaar is op concrete problemen uit de praktijk.

Ook in 2002 was de studiegroep weer een groot succes. U kunt zich hiervan overtuigen door de bijdragen te lezen in deze proceedings. Namens de programmacommissie van Wiskunde Toegepast wil ik mijn dank uitspreken aan de organisatoren en aan iedereen die actief heeft deelgenomen aan de studiegroep. Ik ben ervan overtuigd dat de uitkomsten van deze studiegroep zullen bijdragen aan het beeld van de wiskunde dat wij allen willen uitdragen: niet een stoffig middelbare-schoolvak, maar een uitdagend vakgebied, vooral ook voor aankomende studenten. Die hebben wij nodig om de wiskundigen op te kunnen leiden waaraan onze samenleving dringend behoefte heeft.

Jos de Smit,
Voorzitter van de Programmacommissie
Wiskunde Toegepast

Contents

pp. 1 – 6

Summaries

For the 42nd European Study Group with Industry six problems were selected. The summaries of these problems are presented here.

1. Artis – cooling overheated fish

The Artis Zoo has a problem in its Aquarium and in the adjacent Zoological Museum.

Situation:
Part of the Aquarium is a corridor which contains so called mammoth tanks which measure 5 by 2.5 by 20 meters and are filled with water. Because of the tropical fish inside, the water should have a temperature of 24 degrees Celsius. As there is not much daylight, a dozen lamps have been placed just above the aquarium to make sure the fish inside are visible. However, these big lamps produce a lot of heat.

Problem:
When in summer time the outside temperature reaches 25 degrees, the temperature in the corridor containing the mammoth tanks increases up to 30 degrees. The water itself becomes 27 degrees, which is too much for the fish inside. In the neighbourhood of the lamps, the temperature rises to 40 degrees. Just under the roof sometimes temperatures of 60 degrees have been measured. The museum, adjacent to the aquarium, is suffering from the heat as well: a lot of objects (like stuffed animals) are no longer allowed to be displayed. There are some fans, but these can not do the job, especially not if doors are opened to fight the heat.

Question:
How can we change this situation with a minimum of cost and inconvenience for the visitors, employees and fish?

2. Philips Natlab – compression of audio-signals

Philips Natlab is looking for new ways to compress audio-signals.

Situation:
A new method for the digital representation of high quality audio signals has been introduced as an alternative to the widely used 16-bit recording format

used for CD signals. This new method produces 1-bit samples at a rate that typically is 64 times higher as for CD. For CD the samples are generated at a rate of 44.1 kHz.

The new method results in a raw audio data volume which is 4 times as large as for ordinary CD signals. New storage media provide a huge storage capacity, nevertheless it is beneficial to reduce the required storage capacity. Since the new format is intended for high quality audio signals, popular compression techniques that do change the signals, lossy coding, are unacceptable. This opens up a whole new research area of lossless coding of 1-bit audio signals.

Currently, two main methods have been developed for lossless coding of such 1-bit audio streams. The first, low complexity, scheme uses an adaptive prediction table with run-length residual signal coding. The latter, more elaborate, scheme uses linear prediction with arithmetic coding of the residual signal. In combination with buffering techniques, the methods realise typical average coding gains of 1.3 and 2.1, respectively.

Question:
Can we make compression methods that do better?

To evaluate new proposals, a few short excerpts of 1-bit audio signals will be made available, together with the coding gains achieved with the methods mentioned above.

3. Natuur & Techniek – diffusion of euro coins over Europe

On January 2002 twelve European countries have welcomed the euro as their new coin. The euro coins have a national side, which is different for every country. On top of that three mini states San Marino, the Vatican and Monaco have issued coins with their own image. So there are fifteen different euro coins that can be used in every one of those 15 countries. Therefore, unlike in the past, the coins will not be collected and brought back to their home country. The coins will slowly but surely be spreading over the 15 countries. This is the diffusion of the euro, or euro diffusion.

Because the Belgian and Dutch euro coins form only a fraction of the total number of euro coins, it is to be expected that foreign euros will replace most of the native euros. Interesting questions are: how quickly will the foreign euros take the place of the Dutch euros? How many French coins will we find in one year's time in our wallets? But other questions are also possible!

The diffusion of the euro can be studied in two ways, namely both practically and theoretically. The practical side consists of organizing measurements done by school classes and individuals in the Netherlands and Belgium.

For the theoretical side of the problem the science magazine Natuur & Techniek turned to the Study Group. Unlike the usual way, the problem is not fixed; during the week of the Study Group, the participants will be free to raise interesting questions and hopefully answer them as well. Natuur & Techniek is very interested in the discussion on this problem, and will use the results of the Study Group in an article. A preliminary article already appeared in the January 2002 issue.

The eurodiffusion project is an initiative by the Study Group and the science magazine Natuur & Techniek.

4. Phytocare – better advice to rose cultivators

Phytocare is looking for parameters to grow roses.

Situation:
In Agriculture the key point is optimizing the production at limited costs. For decennia already, one has tried for the best, and currently one cannot think of optimizing the production without a computer anymore. In present-day greenhouses the control of the inner climate is fully automated.

The inner climate could be held constant, but for the optimization of the production it is necessary to adapt the inner climate to the conditions outside the greenhouse. The importance of this can be illustrated by the effect of passing showers: if a rose grower does not anticipate with the, possibly sharp, temperature decrement that is due, this could mean a delay of one week for the production. Hence a swift and adequate reaction is of utmost importance.

Many theoretical models have been made that try to connect the climatic conditions to the resulting production of the weed. Unfortunately, growers do not profit much from the insight obtained by the present models. Most studies aim at one particular type of weed, but the characteristics of various types often differ significantly. Perhaps even more important: the characteristics are not constant throughout the year, whereas the present models account for them with fixed parameters.

Problem:
Phytocare thinks of a new approach. The idea is the following. The climate computer applies specified amounts of moist, light, nutrients, etc. to the plants in a greenhouse. At the same time, the inner climate is measured by the same computer: every 5 minutes the computer provides data on a.o. temperature, humidity and luminiscence in the greenhouse. One thus knows the previous living conditions of each plant with a precision of 5 minutes. One can also measure the production per plant per period from the plants themselves: for tomatoes, for example, the total weight of fruit procuced by a plant within a certain period can be counted. For roses, the weed most

of Phytocare's advises deal with, the production can be measured by the growth of branches per week; indeed, a branch can be harvest as soon as it reaches the required length to be sold. Measuring the production can only be done on a longer timescale. Usually the production is measured every week.

Using the measured climatic conditions in the greenhouse and the production per week, Phytocare would like to find species-specific parameters for the plants, for example by fitting them to the data. As explained before, one of the complications is the fact that the climate measurements are carried out every few minutes, whereas the production can be measured on a weekly basis only. With the parameter values obtained, the approach can be reversed again: are their rules of thumb to be given to the growers, by which their climate control can increase the production at reasonable costs? Different scenarios for advice could be calculated.

Questions:
Can the Study Group make a general model in order to facilitate Phytocare's advice to growers? Can growers be advised how to optimize their production using this model, fitted to the individual grower and his greenhouse by the above described approach? Or will Phytocare at least be able to find out under which conditions the photo-synthetic process of the plants is optimal (optimal production is closely related to optimal photo-synthesis)? Is the model accurate enough in order to calculate whether, for example, certain investments in the greenhouse (to optimize the production) will increase the production sufficiently to justify the investments?

5. Magma Design Automation – component placement on chips

The 'holey cheese' problem

One of the steps in the design process of chips is the positioning of every single component or 'cell' on the chip. The cells are mutually connected by wires. The wiring scheme is given, and in this phase of the design process the positioning of the various cells must be determined. Some (relatively few) cells have a prescribed position.

For the classical positioning problem one considers the chip as a two-dimensional plane; the cells are modelled by rectangles of various sizes. The positioning has to satisfy some conditions:

1 the cells must be placed within a certain rectangle (core area),
2 cells are not allowed to overlap,
3 the total wire length must be minimized.

For this problem many algorithmes are known, each one with its specific pros and cons.

The problem becomes more difficult when large parts of the core area are excluded from positioning, often caused by large, functional components that were placed beforehand (one could think of memory, or components that are designed by other companies). The remaining 'free area' within the core area is usually comparable to a cheese with holes, or 'holey cheese'. Obviously, the cells cannot be placed on the blockages, and this additional requirement makes the positioning problem significantly harder. For a typical 'holey cheese' the free area is strongly disconnected.

The current algorithmes of Magma DA suffice for more or less convex areas. However, in 'holey cheeses' they often end up in local minima that are far from optimal.

Question:
How can Magma DA find good solutions in case of holey cheeses?

6. NIOZ – reconstruction of sea-surface temperatures using fossil marine plankton

The predictive quality of climate models can be enhanced by incorporating information about temperatures from the past. A number of methods have been developed to determine the ancient temperatures of the upper ocean, and one of these is based on the use of deep sea micro fossils.

For many millions of years a large number of species of the invertebrate group planktic foraminifera have lived in the upper water level of the world's oceans. These organisms produce little shells of calcite (CaCO3) that function as a skeleton. Without changing the isotope ratio of oxygen in the dissolved CO2 in the ambient water, the water temperature has a direct influence on the isotope composition in the calcite shells of the plankton. If one would know the isotope composition of the ocean water, one could hence deduct the ocean water temperature from the isotope composition in the calcite shells. The isotope composition of the ocean water from ancient times is practically unknown, however, and, for theoretical reasons, it's not advisable to try to model it either. One way to get around this problem is to take more than one species of plankton:

The different species of plankton don't prefer the same ecological conditions. Some are adapted to live under colder conditions than others. One may hence in principle infer absolute temperature differences from plankton that has lived during the same time-span in the same water level in the same region: the isotope ratio of the water remains fairly constant during relativey short time-spans, but temperatures differ considerably both regionally and in time. Information could thus be obtained about both average temperatures in a given era, and of the variability of these temperatures. The variability is quite large, and therefore interesting to know. Other methods to determine ancient sea water temperatures only obtain mean temperaturs, but the NIOZ is interested to obtain variances as well.

Results obtained so far indicate that the isotope composition in the calcite shells are not only determined by the temperature of the water in which the plankton lived, but also by other ecological influences. For instance, food availability also has its influences on the relative abundance of the different species of foraminifera. It is hence difficult to produce direct conclusions from the isotope composition data from the fossil calcite shells.

Questions:
The challenge now is to construct a model that encompasses the ecology of the organisms, that can still be used to infer the temperatures of the past from the isotope composition data found in the deep sea sedimentary record. It will be very easy to make this model extremely complicated, considering the number of side effects involved. How can we reduce this to a reasonable model that still simulates enough of the observed phenomena?

Participants

Andrei Abramian
Technische Universiteit Delft
andrei@dv.twi.tudelft.nl

Alexei Beliaev
Vrije Universiteit Amsterdam
beliaev@cs.vu.nl

Jan Bouwe van den Berg
University of Nottingham
jan.bouwe@nottingham.ac.uk

Franziska Bittner
Universiteit Utrecht
bittner@math.uu.nl

Piet van Blokland
Hogeschool Holland
pjvanblokland@chello.nl

Onno Bokhove
Universiteit Twente
o.bokhove@math.utwente.nl

Lorna Booth
Universiteit Utrecht
booth@math.uu.nl

Rachel Brouwer
CWI
rachel.brouwer@cwi.nl

Thijs Brouwer
Universiteit van Amsterdam
thijs.brouwer@student.uva.nl

Fons Bruekers
Philips Research
fons.bruekers@philips.com

Chris Budd
University of Bath
cjb@maths.bath.ac.uk

Adriaan van der Burgh
Technische Universiteit Delft
a.h.p.vanderburgh@its.tudelft.nl

Edi Cahyono
Universiteit Twente
e.cahyono@math.utwente.nl

Carlota Cuesta
Vrije Universiteit Amsterdam
carlota@cs.vu.nl

Natalia Davydova
Universiteit Utrecht
davydova@math.uu.nl

Dee Denteneer
Philips Research
dee.denteneer@philips.com

Arjen Doelman
Universiteit van Amsterdam
doelman@science.uva.nl

Johan Dubbeldam
Technische Universiteit Eindhoven
j.l.a.dubbeldam@tue.nl

Koen van Eijk
Magma Design Automation
koen@Magma-DA.com

Barbera van de Fliert
bvandefliert@yahoo.com

Philipp Getto
Universiteit Utrecht
getto@math.uu.nl

Dion Gijswijt
Universiteit van Amsterdam
gijswijt@science.uva.nl

Geertje Hek
Universiteit van Amsterdam
ghek@science.uva.nl

Piet Hemker
CWI
pieth@cwi.nl

Kirankumar Hiremath
Universiteit Twente
k.r.hiremath@math.utwente.nl

Michiel Hochstenbach
Universiteit Utrecht
hochsten@math.uu.nl

Bas van 't Hof
VRtech Computing
bas@vortech.nl

Ale Jan Homburg
Universiteit van Amsterdam
alejan@science.uva.nl

Ed Huijbregts
Magma Design Automation
ed@Magma-DA.com

Joost Hulshof
Vrije Universiteit Amsterdam
jhulshof@cs.vu.nl

Cor Hurkens
Technische Universiteit Eindhoven
wscor@win.tue.nl

David Iron
Universiteit van Amsterdam
diron@science.uva.nl

Lute Kamstra
CWI
lute.kamstra@cwi.nl

Yaroslav Kondratyuk
Universiteit Utrecht
kondratyuk@math.uu.nl

Ger Koole
Vrije Universiteit Amsterdam
koole@cs.vu.nl

Simon Kronemeijer
Universiteit van Amsterdam
kronemj@science.uva.nl

Andreas Kyprianou
Universiteit Utrecht
kypriano@math.uu.nl

Andre Leger
University of Bath
mapajpl@bath.ac.uk

Mervyn Lewis
CWI
mervyn.lewis@cwi.nl

Martijn van Manen
Universiteit Utrecht
manen@math.uu.nl

Jaap Molenaar
TUE/UT
j.molenaar1@tue.nl

Carolynne Montijn
CWI
carolynne.montijn@cwi.nl

Daniel Nitzpon
Universiteit van Amsterdam
nitzpon@gmx.net

Misja Nuyens
Universiteit van Amsterdam
mnuyens@science.uva.nl

Simona Orzan
CWI
simona@cwi.nl

Nick Ovenden
Technische Universiteit Eindhoven
n.c.ovenden@tue.nl

Jun Pang
CWI
jun.pang@cwi.nl

Simon van der Pal
Artis
vanderpal@artis.nl

Frank Peeters
NIOZ
peeters@nioz.nl

Mark Peletier
CWI
peletier@cwi.nl

Gemma Piella
CWI
piella@cwi.nl

Derk Pik
Universiteit Leiden
drpik@math.leidenuniv.nl

Bob Planqué
CWI
rplanque@cwi.nl

Harmen van der Ploeg
Universiteit van Amsterdam
hvdploeg@science.uva.nl

Iuliu Sorin Pop
Technische Universiteit Eindhoven
i.pop@tue.nl

Georg Prokert
Technische Universiteit Eindhoven
g.prokert@tue.nl

Marieke Quant
Katholieke Universiteit Brabant
quant@kub.nl

Vivi Rottschäfer
Universiteit Leiden
vivi@math.leidenuniv.nl

Jacques Rougemont
Heriot-Watt University
j.rougemont@ma.hw.ac.uk

Dick van der Sar
Phytocare
dick.vdsar@phytocare.nl

Jan A.M. Schreuder
Universiteit Twente
j.a.m.Schreuder@math.utwente.nl

Piet Sondervan
Artis

Bernadetta Tarigan
CWI
bernadetta.tarigan@cwi.nl

Guido Terra
UvA/NIOZ
gmterra@science.uva.nl

Paul Dario Toasa
University of Kaiserslautern
toasa@mathematik.uni-kl.de

Evgeny Verbitskiy
Eurandom
verbitskiy@eurandom.tue.nl

Erik Vermeulen
Natuur & Techniek
evermeulen@natutech.nl

JF Williams
University of Bath
j.f.williams@maths.bath.ac.uk

Djoko Wirosoetisno
Universiteit Twente
djoko@maths.ed.ac.uk

Dmitri Znamenski
Vrije Universiteit Amsterdam
dznamen@cs.vu.nl

pp. 41 – 57

CHAPTER 3

The Euro Diffusion Project

Piet van Blokland, Lorna Booth, Kirankumar Hiremath, , Michiel Hochstenbach, Ger Koole, Sorin Pop, Marieke Quant, Djoko Wirosoetisno.

ABSTRACT. From 1st January 2002 we have the unique possibility to follow the spread of national euro coins over the different European countries. We model and analyse this movement and estimate the time it will take before on average half the coins in our wallet will be foreign.

KEYWORDS: Markov chains, diffusion

1. Introduction

On January 1, 2002 a total of 12 European countries replaced their national currencies with the euro. These coins are not identical, one side of each of the 8 denominations of coins differs from country to country, and as people travel these coins mix. The national banks have decided that no redistribution will take place, and therefore it is expected that in the long run a close to perfect mixture of coins will take place. In this paper we analyse this problem and come up with a mathematical model for the movement of coins that enables us to estimate the speed at which this process will occur. This speed strongly depends on the value of certain parameters that are hard to estimate and for this reason we have to be careful with conclusions. However we estimate that it will take around 12 months before roughly half of all coins in Dutch wallets will be foreign.

Our other main conclusions are:

- Markov chains are the obvious models to model the movement of euro coins,
- Continuous models are the natural approximation of these,
- The data on the Euro Diffusion web site are unreliable and must be used for parameter estimations only with care.

In this paper each section or subsection is marked with a number of stars to indicate approximate level of difficulty:

* means that the section can be understood without effort,

** means that the section contains mathematics which is explained very gently,

*** means that the section would be accessible to anyone who has done a little university level mathematics, or who is prepared to work a bit more at understanding.

2. Basic facts*

The euro was introduced in 12 European countries at the beginning of 2002. A total of 64.9 gigacoins (1 GC = 1 000 000 000 coins) were made, of which 3.3 GC are Dutch, but not all these coins were brought into circulation right away: in Holland 1.6 GC were put into the market on 1st January 2002. The number of coins made by each country is not in proportion with the population: e.g., France made around 190 coins per person, Germany 280, and Holland 200. The reason for these differences is unknown and is part of the policy of each national bank.

New coins are still being brought into circulation after the introduction of the euro, mainly to compensate for savings and wastage. At the end of 2001 all money-boxes were emptied of their national currencies, and now they are slowly being filled again with euros. The Dutch national bank (DNB) expects that this will require about 100 MC per month in the Netherlands. The amount of coins saved is eventually expected to equal the amount in active circulation; in the guilder age it was estimated that 1.5 GC out of a total of 3.0 GC were saved. Wastage is considerably less: in the guilder age it was about 50 to 100 MC per year, largely due to loss (people dropping coins, etc.) and coins going abroad forever. We think that wastage will be a little less as tourists from other countries can also use the Dutch euro coins in other European countries, and therefore 60 MC, 5 MC per month, seems to be a reasonable estimate of this type of loss. The fact that there are new coins brought into circulation will mean that the mix will never be perfect, but as the waste is tiny compared to the total number of coins in circulation, this effect will be small.

A new effect, which is hard to estimate, is the collection of euro coins. Collection used to be an effect of negligible size; but nowadays many people seem to collect foreign euro coins. The size of this effect and when these coins will be brought into active circulation again are difficult to estimate.

When counting the number of coins in our wallets, we usually find 10 to 20 coins, and this is confirmed by a study from DNB which shows that we carry on average 15 coins. However, per person 100 coins were brought into circulation. The remaining 85% of all coins can be found at check-outs in shops and at banks. This may imply that movement of the euro coins will be relatively slow: only 15% of all coins can be taken abroad at any one time! For example, suppose that 2/3 of the Dutch population go abroad during the summer holidays, and assume that after their foreign visits their coins are representative of the coin mixture in the country they visited (which will certainly contain a (small) percentage Dutch coins!). This replaces only 10% of all Dutch coins with (mainly) foreign coins.

People travelling and taking their coins abroad is not the only possible reason for the spread of euro coins. Another is the possibility that Dutch banks will buy euro coins from somewhere other than DNB, which might be cheaper, for example, if a depot of the Belgian or German national bank is closer to the particular Dutch bank. The effects of this and how national banks will react to it are difficult to predict; nobody knows if it will occur and to which extent. This effect might even create an imbalance between the quantities of coins in a country: there might be a net flow of coins in or out a country.

3. Data

3.1. Observations*. To fit the parameters in our model, we use the data gathered in the web-site
`http://www.wiskgenoot.nl/eurodiffusie/`
which is built up by voluntary observations, mainly from different parts of the Netherlands and Flemish Belgium, and from a few points elsewhere. Each data point consists of the date and locale, and the number of coins the correspondent has, along with the breakdown according to denominations and countries of origin. Various summaries are available on the web-site.

We summarise the observations for the percentage of 1-euro coins in the Netherlands in Table 1.

t	nl	be	de	fi	fr	gr	ie	it	lu	os	pt	es	#no.
1	90.9	1.1	4.2	0.0	1.1	0.0	0.0	1.1	0.3	0.0	0.3	0.8	353
2	84.7	3.6	4.4	0.0	2.9	1.7	0.4	0.1	0.3	1.0	0.0	0.8	724
3	82.7	5.2	5.1	0.2	3.6	0.0	0.1	1.0	0.7	0.6	0.0	0.8	1186
4	85.2	2.8	4.7	0.2	2.3	0.2	0.2	1.1	0.1	2.2	0.2	0.6	3524
5	80.6	5.5	5.7	0.2	4.3	0.0	0.0	1.0	0.1	1.3	0.2	0.9	976
6	75.3	5.8	8.9	0.1	3.3	0.2	0.5	1.6	0.0	2.5	0.2	1.7	1260
7	79.4	5.4	6.9	0.2	2.5	0.1	0.4	1.2	0.4	2.2	0.2	1.2	3977
8	74.7	10.3	5.4	0.1	3.7	0.3	0.1	1.1	0.5	1.9	0.6	1.3	874
9	80.0	5.6	7.1	0.4	1.9	0.6	0.4	1.2	0.0	1.3	0.4	1.2	520

TABLE 1. Percentage of 1-euro coins of various types in the Netherlands

Here line `1` corresponds to an average over the first ten days in January and so on; the countries are labelled by their internet abbreviations; the last column is the total number of coins counted, to provide some indications of the error. Similarly we see the number of 1-euro coins for Belgium in Table 2.

In Figure 1, we plot the percentages of local coins in the Netherlands and Belgium, with the horizontal axis being the time in days.

t	nl	be	de	fi	fr	gr	ie	it	lu	os	pt	es	#no.
1	0.0	94.6	2.7	0.0	2.7	0.0	0.0	0.0	0.0	0.0	0.0	0.0	37
2	5.3	89.8	1.9	0.0	1.9	0.0	0.0	0.3	0.6	0.0	0.0	0.3	323
3	5.1	87.6	1.8	0.1	3.3	0.1	0.0	0.6	0.3	0.5	0.0	0.4	930
4	5.4	85.7	2.5	0.1	3.4	0.0	0.0	0.6	0.6	0.8	0.1	0.6	4045
5	5.1	74.2	4.5	0.3	10.3	0.1	0.3	1.2	1.5	1.4	0.0	1.3	1027
6	6.7	74.4	4.9	0.1	7.3	0.0	0.2	1.5	1.0	2.5	0.2	1.2	2275
7	6.8	75.9	4.4	0.1	6.3	0.2	0.2	1.7	1.8	1.3	0.1	1.2	4263
8	5.7	82.6	3.0	0.4	3.1	0.1	0.4	1.2	0.6	0.4	0.1	2.3	1376
9	5.2	77.2	3.2	0.0	5.5	0.0	0.8	1.5	1.2	0.2	0.0	5.0	400

TABLE 2. Percentage of 1-euro coins of various types in Belgium

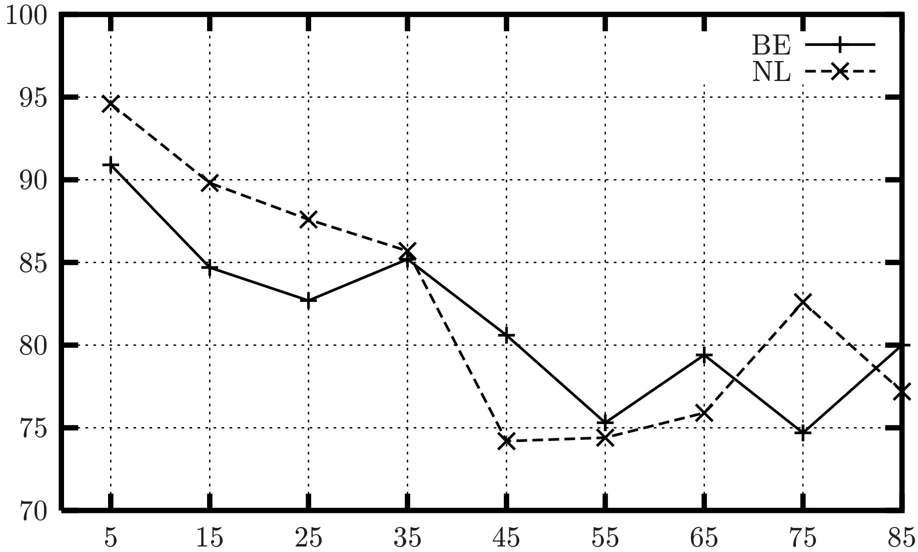

FIGURE 1. Percentage of local coins in the Netherlands and Belgium over the first 85 days

Although a general downward trend can be seen for both curves, the size of the fluctuations is rather puzzling, considering the fact that the numbers of coins counted are quite large (corresponding to hundreds or thousands of people). A similar irregular behaviour is also observed when the data is restricted to smaller areas (e.g., greater Utrecht or Brussels).

We make a couple of observations about the data:

0. Not surprisingly, the percentage of local coins at any given location tend to decrease over time. On February 1, these numbers are 91% for the Netherlands and 89% for Belgium (averaged over all denominations); the corresponding figures are 87% and 81%, respectively, for March 1.
A possibly more illuminating example is the case of Luxembourg: *72% of

coins are (still) of local origin on February 1, and *68% on March 1 (data is from the Euro Diffusion web-site, asterisks denote possible inaccuracies.) This suggests that the mixing process is likely to take several years.

1. The observed data for the Netherlands and Belgium tell us that coins of different denominations mix at different rates. At any given time, the percentage of foreign coins is higher the higher the denomination is: In the Netherlands on March 1, only 8.3% of the 1-cent coins are of foreign origin; for 10c it is 11.7%, and for 2-euro it is 19.6%.

2. Data for the Netherlands and Belgium show apparently anomalously large proportions of Austrian coins, being 2.3% and 1.7% on March 1 (well on its way to the long-time limit of 3.1%). Data from the Canary Islands (Spain) and Creta (Greece) show that foreign coins are almost exclusively German. These suggest that there are considerable variations in the travel habits of people from different countries; one may attempt to determine certain aspects of these in addition to the questions posed above.

The fluctuation in the data may be partially explained by the fact that at the Euro Diffusion web site everybody on the Internet can enter measurements of the number of different euro coins in their wallets. For this reason there is no guarantee that these measurements represent the real situation: the selection of people is probably biased and the moments at which they enter data might be strongly biased (e.g., somebody enters data when he has something to tell, i.e., he has some "rare" coins in his wallet!). The extent to which these biases corrupt the data is difficult to measure without further study, so to get some feeling for it we did some counts ourselves during the workshop. The results can be found in Table 3.

Who/where	date	NL	B	other countries
Lunch (Amsterdam)	19 Feb	293/93.1%	12/3.8%	10/3.1%
Lorna's Students (Utrecht)	20 Feb	125/98.4%	2/1.6%	0/0.0%
Check-out K.d.V. Inst. (A'dam)	20 Feb	233/94.7%	3/1.2%	10/4.1%

TABLE 3. Measurements during the SWI workshop in February 2002

As an alternative we tried to get the, in our opinion, most reliable data from the web site. We looked for unique "eurometers", with multiple participants, that have entered data three times. The results can be found in Table 4.

Eurometer	dates	percentages NL
46 (Rotterdam)	18/1, 31/1, 8/2	98, 96, 98
167 (Utrecht)	21/1, 11/2, 18/2	95, 98, 96
510 (Hengelo (O))	15/1, 24/1, 1/2	96, 92, 91

TABLE 4. Selected measurements from Euro Diffusion database

It is surprising to see that, also in the measurements considered to be reliable, there is a high variability in outcome, despite the size of the measurements. This suggests that the measurements come from different distributions.

4. A First Estimate**

To come to an estimate of the whereabouts of coins on 1st February we assume that diffusion is approximately linear during January and February. Thus the measurements at Feb 19 and 20 of Table 3 should be multiplied by roughly 3/5 to find an estimate for the end of January. For the measurements in Table 4 similar factors are used. This gives the following estimate (numbers correspond to own measurements and "reliable" eurometers, respectively):

$$\text{(1)} \quad \Big[\frac{3}{5}(6.9+1.6+5.3)+\frac{3}{2}2+4+\frac{3}{4}2+\frac{3}{2}5+\frac{3}{4}2+ \frac{3}{5}4+2\times 4+\frac{5}{4}8+9\Big]/12 \approx 4.6$$

This result is also unreliable, and biased by the high number of non-Dutch euros measured by eurometer 510, therefore as a first estimate we took 4% diffusion per month. This has to be verified during later months, although even the March 1 measurement will be disturbed by the influence of the February holidays.

Note that we assumed that the percentage of foreign coins equals the rate at which the coins move between the Netherlands and abroad, which is a good approximation as long as the diffusion is slow and we are near to the beginning of the process. In case of an increased rate this is no longer the case, then we have to use more sophisticated models, and estimate based upon these models. In the next section we will describe the types of models we use, and after that we will explore a number of techniques for estimating the parameters more accurately.

5. Modelling the Spread of Euros with Markov Chains

Following the grand total of 64.9 billion euro coins in some mathematical model is impossible. It is also useless: the movement of a single coin shows the same (type of) behaviour as any other coin, although this may depend on the country of origin and denomination. Therefore we may concentrate on a single coin and follow it in a probabilistic manner its way around Europe. From this we can then draw conclusions for the entire population of euro coins.

The movement of some arbitrary coin can be followed by a *Markov chain*. The simplest Markov chain model has only two states but can already be used to give a first simple model for the movement of coins. We will describe this model in detail, discuss its shortcomings, and deal in less detail with more sophisticated models.

5.1. Two-state model.** In the simplest model we have two states: a coin can either be in Holland or abroad. We denote these states with H and A. We neglect wastage, and do not distinguish between coins in active circulation and saved coins. We also assume that the net rate of coins into the Netherlands is equal to the net rate out. Let d be the rate at which diffusion occurs: this means the percentage of coins in Holland that is replaced by coins from abroad after one month. Note that it is *not* the rate at which non-Dutch coins arrive: foreign coins in Holland may go abroad again, and Dutch coins abroad may be taken back to Holland! As we can assume that the net flow equals 0, we know that (on average) every coin going abroad is replaced by a coin coming from abroad. Thus having d percent of foreign coins in the Netherlands after 1 month corresponds to there being a probability of d that an arbitrary coin "goes" abroad and doesn't return in one month. Thus p_{HA}, the probability of a transition from H to A in the 2-state Markov chain, is d and, assuming an active circulation of 1.6 GC in the Netherlands, the flow out in one month is $1.6d$ GC. As the flow in equal the flow out, we find $p_{AH} = \frac{1.6d}{ACA}$, with ACA the active circulation abroad. Assuming that the active circulation abroad is the proportion of coins in active circulation in the Netherlands multiplied by the total number of coins abroad, we see that $ACA = \frac{1.6}{3.3}61.6 = 29.9$ GC. Using this we find that if $d = 0.04$ (see section 3) then $p_{AH} = 0.00214$. Now that we know the transition probabilities we can start our computation. Let Q be the *transition matrix*:

$$Q = \begin{pmatrix} 1 - p_{HA} & p_{HA} \\ p_{AH} & 1 - p_{AH} \end{pmatrix} = \begin{pmatrix} 0.96 & 0.04 \\ 0.00214 & 0.99786 \end{pmatrix}.$$

The matrix Q should be interpreted as follows: The first row shows where a coin, initially in Holland, will be after 1 month. The first entry shows the probability that it will be in Holland, the second entry that it will be abroad. The second row corresponds to a coin originating from abroad. Again, the number on the diagonal, Q_{22}, is the probability that the coin is abroad after 1 month; the left hand element is the probability that the coin is in Holland. Multiplying Q with itself gives numbers with the same interpretation, but for a 2-month period, and, similarly, Q^n gives the transition probabilities for n months. For example lets look at what happens after 15 months. Computation, with $d = 0.04$, shows that

$$Q^{15} = \begin{pmatrix} 0.54841 & 0.45159 \\ 0.02639 & 0.97361 \end{pmatrix}.$$

Thus, 15 months with a diffusion at the rate of January will lead to 45% of Dutch coins being abroad. They are replaced by foreign coins, and thus there are 45% foreign coins in the Netherlands. coins.)

We expect that the speed of diffusion is temporarily increased by the summer holidays and the February skiing holidays. Assume that 2 out 3 people go on holidays during summer to another country which uses the euro. As 15% of all coin are in our wallets, this will result in 10% diffusion. At a 4% diffusion rate, this is the equivalent of 2.5 months. Together with the effect of the skiing holidays we assume that the combined holidays count for 3 to 4 months. Thus the situation by the end of the year is equivalent to 15 to 16 months at rate 0.04. To conclude, we expect that the percentage of foreign coins will be 50% somewhere in the first months of 2003.

5.2. Savings*. We can model the effects due to people saving in a similar way. We now have another state - the piggy-bank - from which coins are released only after a long time. To compensate for this the Dutch National Bank puts new (Dutch) euros into circulation until the piggy-banks are full. Simulations show that this slows the diffusion of coins, for two reasons. The first is the addition of Dutch euros and the second is that the piggy-banks act as a reserve of Dutch euros which are released later into the money flow.

5.3. Refinements*. The simple model has $(0.05085, 0.94915)$ as equilibrium, which is is not realistic as wastage is ignored. Wastage leads to shortages of coins that lead to the production of new coins. These new coins are then (presumably) released in the home country, leading to a slightly higher percentage of national coins in each country.

In the simple model all countries are aggregated in a single state, however it is to be expected that diffusion occurs faster with certain countries than with other countries. A model with multiple states representing different countries could make predictions on the diffusion of other coins.

Finally, a regional approach, splitting countries up in regions, can be followed. This would be of particular value if we want to use the data from the Euro Diffusion web-site to the full. In this case we could use a model that has all of the regions (or groups of the regions) used by the web-site, to estimate how quickly coins move around inside the Netherlands and Belgium.

As also mentioned in Section 5.1, seasonal variations will have non-negligible effects on the transition to an equilibrium. It may be almost impossible to insert these variations into the model in a realistic way.

6. Markov Chains and Continuous Models**

In the last section we used Markov Chains to model the movement of coins around. Each coin jumped around independently with the probabilities given in the matrix Q. For example, the chance that a coin that was in the Netherlands moves outside it after one month is 0.04. However if we have a very large number of coins, as we do, roughly 4 per cent of the coins that are the Netherlands will move outside of it. Although each coin chooses to move or not independently of the other coins, there is a mathematical theorem, the Law of Large Numbers, which tells us that the actual proportion of coins moving will be very close to 4 per cent. Therefore if we assume a continuous model, in which *exactly* 4 per cent of the coins move from the Netherlands to somewhere else, we will be close to the truth. In fact the error in the model will be roughly a constant divided by the square root of the number of coins, which really is very small. Therefore in Sections 7 and 8 we consider continuous models.

7. Parameter estimation based on data from multiple months**

Suppose we have a matrix Q, like that in Section 5.1, which tells us how coins move about. The entry in the ith row and jth column tells us what proportion of coins from region i move to region j in one month. In the example in section 5.1 we had just two regions, the Netherlands (region 1) and the rest of the world (region 2), but we could instead have a model with more regions. If we wanted to make the most possible use of the data from the Euro Diffusion web-site we might end up having 70 regions, more than 60 parts of the Netherlands and Belgium and the other 10 countries. We might also want to have regions representing areas outside of Europe, or the piggy-banks (saving pots) of people in different countries, or even the place where all of the lost money goes. Suppose we have r of these regions and give each region a number from 1 to r. This makes Q an $r \times r$ matrix, with r^2 entries, none of which we know. We want to estimate Q using the data we have from the web-site.

Suppose that we consider m types of coin. These could be the 12 nationalities, in which case m would be twelve, or we could decide to group some nationalities together as in Section 5.1 in which we had two types, Dutch coins (type 1) and all others (type 2). We label each type with a number from 1 to m. We know or can find out in which region(s) each type of coin begins, on 1st January 2001. For example in Section 5.1 we knew that all of the Dutch coins began in the Netherlands while all of the non-Dutch coins began in the rest of the world. So for each type of coin we can make a vector, b, which tells us where those coins are. In the example we have $b(1) = (1, 0)$ which says that of the 1st type of coins (the Dutch ones) all are in region

1(the Netherlands) and none are in region 2 (the rest of the world), and $b(2) = (0, 1)$ which says that of the 2nd type of coins (the non-Dutch ones) none are in region 1 (the Netherlands) and all are in region 2 (the rest of the world). We then have m of these vectors $b(1), b(2), b(3), \ldots, b(m)$ (one for each type of coin) and each will have r entries.

We also have some measurements from the web-site. We call the number of measurements of coin type c in region i after t months $n(c, i, t)$. Suppose we measure the number of coins of each type in each region every month for s months. In the next sections we will discuss a number of possible estimation techniques.

Unfortunately all of these estimates inevitably run the risk of significant errors, due to the issues discussed in Section 3.

7.1. Maximum Likelihood Estimation*.** A common method in this type of problem is called *maximum likelihood estimation*. To use this we first write down the probability that we would see the measurements we have seen if we knew what Q was. If we had samples from all of the regions this turns out to be

$$p(Q) = \prod_{t=1}^{s} \prod_{i=1}^{r} N(i,t) \prod_{c=1}^{m} ((b(c)Q^t)_i)^{n(c,i,t)}$$

where

$$N(i,t) = \binom{\sum_{c=1}^{m} n(c,i,t)}{n(1,i,t)n(2,i,t)\cdots n(m,i,t)},$$

a multinomial coefficient. What we then do is find the matrix Q which maximises this probability.

This does look horrible, but fortunately we don't have to worry about $N(i,t)$. It is possible to show that the value of Q that maximises this probability also maximises:

$$\hat{p}(Q) = \prod_{t=1}^{s} \prod_{c=1}^{m} \prod_{i=1}^{r} ((b(c)Q^t)_i)^{n(c,i,t)}.$$

There are two things we do have to keep in mind when we do this. The entries of Q are proportions and are therefore all larger than or equal to zero. The entry in the ith row and jth column tells us what proportion of coins from region i move to region j, so if we sum all the entries is a row we should get the proportion of coins that go anywhere or stay in the same country, i.e. 1.

In formulas we need to

$$\text{maximise} \prod_{t=1}^{s} \prod_{c=1}^{m} \prod_{i=1}^{r} ((b(c)Q^t)_i)^{n(c,i,t)}$$

$$\text{subject to } Q(i,j) \geq 0, \text{for all } i \text{ and } j, \text{ and } \sum_{j=1}^{r} Q(i,j) = 1 \text{ for all } i$$

Generally this is a horrible problem which isn't easy to solve using mathematics and so we will turn to computers to do this numerically. This will typically be hard and slow.

7.2. An Iterative Approach***.** The "true" proportion of coins of type c in region i after t months is $(b(c)Q^t)_i$. A logical estimate for this is

$$p_t(c,i) = n(c,i,t) / \sum_{c'=1}^{m} n(c',i,t), \tag{2}$$

i.e. the proportion of coins of type c that we see in region i after t months.

One of the problems that we have with the estimation is that we do not have measurements from every region. If we did we could estimate the row vector $b(c)Q^t$ by the row vector $(p_t(c,1)\ p_t(c,2)\ p_t(c,3)\ \ldots\ p_t(c,r))$. Notice that if we stack up the row vectors $b(c)Q^t$ we get,

$$\begin{pmatrix} b(1)Q^t \\ b(2)Q^t \\ \vdots \\ b(m)Q^t \end{pmatrix} = \begin{pmatrix} 1 & 0 & 0 & \cdots & 0 \\ 0 & 1 & 0 & \cdots & 0 \\ 0 & 0 & 1 & \cdots & 0 \\ \ldots & \ldots & \ldots & \ddots & \ldots \\ 0 & 0 & 0 & \cdots & 1 \end{pmatrix} Q^t = Q^t$$

Therefore the matrix,

$$M_t = \begin{pmatrix} p_t(1,1) & p_t(1,2) & p_t(1,3) & \cdots & p_t(1,r) \\ p_t(2,1) & p_t(2,2) & p_t(2,3) & \cdots & p_t(2,r) \\ \ldots & \ldots & \ldots & \ddots & \ldots \\ p_t(m,1) & p_t(m,2) & p_t(m,3) & \cdots & p_t(m,r) \end{pmatrix}$$

this would give us a good estimate for Q^t, and we could find an estimate for Q by taking this matrix to the power of $(1/t)$. We could then combine the estimates from each timestep to give a final estimate for Q,

$$\frac{1}{s}\sum_{t=1}^{s} M_t^{1/t}.$$

One way we can deal with this missing data is to use an iterative algorithm. We begin with a guess for Q, call it Q_0, and gradually improve this. We begin with j (a counter) set to 0.

Step One: Let $p_t(c,i) = (b(c)Q_j^t)_i$ for all coin types and each of the regions i that we do not have data from.

Step Two: For each t, form M_t using Equation 2 for the regions we have data from and the estimates from Step One for the regions where we do not have data.

Step Three: Increase j by 1, let $Q_j = \frac{1}{s}\sum_{t=1}^{s} M_t^{1/t}$, and go to Step One.

This should converge to a reasonable estimate for Q, in a reasonable amount of time.

7.3. An Example*.** To illustrate the methods above we will apply them to a three state model, in which we take the Netherlands (region 1), Belgium (region 2) and everywhere else (region 3) as our regions. We also consider three types of euro coin, Dutch (type 1), Belgian (type 2) and the others (type 3). Therefore $b(1) = (1,0,0)$, $b(2) = (0,1,0)$ and $b(3) = (0,0,1)$, as all of the coins begin in their respective countries. For the maximum likelihood we consider two measurements, that of 1st February and that of 1st March. For the iterative approach we also consider 1st April.

The data we have from the Euro Diffusion web-site is that:

$$
\begin{array}{ll}
n(1,1,1) = 0.908 & n(1,2,1) = 0.041 \\
n(2,1,1) = 0.023 & n(2,2,1) = 0.888 \\
n(3,1,1) = 0.069 & n(3,2,1) = 0.071 \\
n(1,1,2) = 0.871 & n(1,2,2) = 0.06 \\
n(2,1,2) = 0.031 & n(2,2,2) = 0.807 \\
n(3,1,2) = 0.098 & n(3,2,2) = 0.133 \\
n(3,1,2) = 0.098 & n(3,2,2) = 0.133 \\
n(1,1,3) = 0.823 & n(1,2,3) = 0.064 \\
n(2,1,3) = 0.045 & n(2,2,3) = 0.807 \\
n(3,1,3) = 0.132 & n(3,2,3) = 0.129
\end{array}
$$

Maximum Likelihood

In this case we don't have measurements for the numbers of coins outside the Netherlands and Belgium. Therefore for the maximum likelihood we need to

$$\text{maximise} \prod_{t=1}^{2}\prod_{c=1}^{3}\prod_{i=1}^{2}((b(c)Q^t)_i)^{n(c,i,t)}$$

$$\text{subject to } Q(i,j) \geq 0, \text{for all } i \text{ and } j, \text{ and } \sum_{j=1}^{r} Q(i,j) = 1 \text{ for all } i$$

If we write this out, filling in the values we know and writing

$$Q = \begin{pmatrix} q_{11} & q_{12} & q_{13} \\ q_{21} & q_{22} & q_{23} \\ q_{31} & q_{32} & q_{33} \end{pmatrix},$$

this becomes the problem of maximising,

$$q_{11}^{0.908} q_{12}^{0.041} q_{21}^{0.023} q_{22}^{0.888} q_{31}^{0.069} (q_1 1^2 + q_{12}q_{21} + q_{13}q_{31})^{1.742} \times$$

$$(q_{12}q_{21}+q_{21}q_{22}+q_{23}q_{31})^{0.031}q_{32}^{0.071}(q_{11}q_{12}+q_{12}q_{22}+q_{13}q_{32})^{0.06}\times$$
$$(q_{12}q_{21}+q_{22}^2+q_{23}q_{32})^{0.807}(q_{11}q_{31}+q_{13}q_{32}+q_{31}q_{33})^{0.098}\times$$
$$(q_{12}q_{31}+q_{22}q_{32}+q_{32}q_{33})^{0.133}$$
$$\text{subject to } q_{ij}\geq 0, \text{for all } i \text{ and } j, \text{ and } \sum_{j=1}^{r} q_{ij}=1 \text{ for all } i$$

This really hard, even for a computer, but a preliminary numerical investigation suggests that this is maximised by a matrix close to:

$$\begin{pmatrix} 0.975 & 0.025 & 0 \\ 0.02 & 0.98 & 0 \\ 0.4 & 0.6 & 0 \end{pmatrix}.$$

The Iterative Approach

The iterative approach is easily programmed and the estimates converge quickly to:

$$\begin{pmatrix} 0.924326 & 0.0172391 & 0.0584345 \\ 0.0309947 & 0.903493 & 0.0655122 \\ 0.0578321 & 0.065527 & 0.876641 \end{pmatrix}.$$

8. A Continuous Model for Euro Movement**

Coins of neighbouring countries mix at boundary cities. Coins of various origins mix with each other at tourist places, at airports, in highway restaurants, gasoline stations etc. This is a typical *continuous* diffusion phenomenon like the mixing of different gases or the spread of smoke or fragrance. The process of continuous diffusion frequently occurs naturally and plays an important role in many applications. The main cause for diffusion is gradient profile, i.e. different concentrations in different places. The coefficient of diffusion (denoted by D) is the measure of the continuous diffusion.

In the long-time limit, and in the absence of sources (which we think is quite a good approximation), it is clear coins of different origins will have roughly the same proportion everywhere within the euro-zone. The natural question to ask is, over what timescale is this going to happen?

In our problem we may be able to use the country-of-origin breakdown of euro coins (which are identical for practical purposes) to deduce information on coin transport (and hence certain aspects of people's economic behaviour) which would be very difficult to measure by other means.

8.1. Model*.** In this section we seek to design a continuous model which, with the least amount of complexity, will best fit the available data (the first three months of 2002) and will make useful predictions of its future behaviour.

As the basis of this model, we shall make the assumption that people are indifferent to origin of coin, i.e., they treat coins of different countries equally. Also, in view of observation **1** in Section 3.1, we shall consider a fixed denomination, say, one euro, in our model.

In this model, we consider two different processes that contribute to the dispersal of coins. The first process is a local one, which arises from people carrying out their daily activities: going to the market, the bank, restaurants, etc. We model this by a diffusion equation, with a diffusion constant $D(x,t)$. The second process is inherently non-local, arising from medium- and long-distance travels. We model this process, a priori, by an integral term. In addition to these, we also need to include sources (bank issues) and losses.

We note that, in reality, there is likely a continuous spectrum of "coin transport scales"; nevertheless, for conceptual purposes and modelling feasibility, we assume that there is a "separation of scale" between local and non-local processes.

Let $m^{(n)}(x,t)$ denote the (normalised; see below) density of 1-euro coins of country n at location x at time t, where $x \in \Omega =$ euro-land and where we take 2002 January 1 to be $t = 0$. A general governing equation for $m^{(n)}(x,t)$ then reads

(3)
$$\frac{\partial}{\partial t} m^{(n)}(x,t) = f^{(n)}(x,t) - \nu(x) \cdot m^{(n)}(x,t) + \nabla \cdot (D(x,t)\nabla m^{(n)}(x,t)) + \int_\Omega K(x,z,t)\, m^{(n)}(z,t)\, \mathrm{d}z.$$

Here the terms are

$f^{(n)}(x,t) =$ source density such as bank,
$\nu(x) =$ rate of loss of coins, assumed independent of t,
$D(x,t) =$ diffusion coefficient,
$K(x,z,t) =$ integral kernel for non-local transport.

We note that if the loss term $\nu(x)$ is independent of x and t, it can be dropped provided that one rescale: $m^{(n)}(x,t) \mapsto \mathrm{e}^{-\nu t} m^{(n)}(x,t)$, this being irrelevant if one only considers the *ratios* (such as percentages) of coins.

Let $e(x)$ denote the population density at point x, assumed to be constant in time. To a very reasonable approximation, we may take the initial conditions to be

$$m^{(n)}(x,0) = \begin{cases} e(x) & x \text{ is in country } n \\ 0 & \text{otherwise.} \end{cases} \tag{4}$$

The possibility (not considered further here) that the coin/population ratio may vary by country can be taken care of by normalising $e(x)$ to account for this factor.

Let us now construct a general model for the integral kernel $K(x,z,t)$. As before, we assume that long-distance travel does not change the population density $e(x)$. Let $p(x,z,t)$ be the probability that an inhabitant of x visits z. Then the integrand $K(x,z,t)m^{(n)}(z,t)$ in (3) can be modelled as,

$$\begin{aligned} K(x,z,t)\,m^{(n)}(z,t) &= \varepsilon_0(x)p(x,z,t)e(x)[m^{(n)}(z,t) - m^{(n)}(x,t)] \\ &\quad + \varepsilon_0(z)p(z,x,t)e(z)[m^{(n)}(z,t) - m^{(n)}(x,t)]. \end{aligned} \tag{5}$$

Here the first term corresponds to residents of x travelling to z and the second to residents of z travelling to x; ε_0 is a numerical coefficient related to the number of coins people carry on long-distance trips. To simplify this further, we may assume (somewhat unrealistically) that "tourists" of different origins are distributed equally among all popular destinations; this would allow us to write $p(z,t)$ on the first line, measuring how popular z is, and $p(x,t)$ on the second.

In what follows, we take $D(x,t)$ to be inversely proportional to the population density $e(x)$, which is effectively constant over the timescales considered here. The reason for this is that people tend to travel farther to the "local shop" in sparsely populated areas than in urban centres, and, assuming that city and rural shops have equal number of customers, this distance is proportional to $e(x)^{-1/2}$. Thus,

$$D(x,t) = D_0/e(x), \tag{6}$$

where the constant D_0 is to be determined from a fit with the observed data.

Instead of the full integral kernel $K(x,z,t)$ in (3), we shall model the non-local effects by "connecting" a small number (10 in this work) of "major airports"; by this we mean to account for all types of long-distance travels, not only by air. We denote each airport by $L_i \subset \Omega$, having a fixed area δ (one element in the numerical model below). For each airport L_i, we assign the number of people served by it, $\rho(L_i)$, and we assume that the number of passengers travelling from airport L_i to L_j is a constant ε' times $\rho(L_i)\rho(L_j)$. The coin transport due to this process can then be modelled by adding the following term to (3) above: For $x^* \in L_i$ and $z^* \in L_j \neq L_i$,

$$\left.\frac{\partial}{\partial t}\right|_{\mathrm{nl}} m^{(n)}(x^*,t) = \frac{\varepsilon}{\delta} \sum_{L_j} \int \rho(x^*)\,\rho(z^*)\left[m^{(n)}(z^*,t) - m^{(n)}(x^*,t)\right] \mathrm{d}z^* \tag{7}$$

where the constant ε (which like D_0 is to be determined from a fit with the data) is related to ε'. The value of δ is irrelevant if it is sufficiently small; one can take the limit $\delta \to 0$ if desired.

Since the population density data is readily accessible (e.g., from a school atlas), our model so far only has two parameters, D_0 and ε, that need to be determined. Noting that we can scale time by D_0, as far as the *qualitative* behaviour of our model so far is concerned, only one parameter matters: the ratio ε/D_0.

The model can also be easily modified (by changing the diffusion constant) to investigate the economic importance of national borders, e.g, whether people living close to the border are more or less likely to cross the border to shop.

8.2. Numerics and Fit.** We ran numerical simulations of a "finite-element" version of the model described above, consisting of 100 cells, each representing 3 million people, which are connected in a rough approximation to the actual geography. The variables are $m_i^{(n)}$, with $n \in \{\text{NL, BE, DE, } \cdots\}$ and $i = 1, \cdots, 100$. The diffusion term

$$\nabla \cdot (D_0 e(x) \nabla m^{(n)}(x,t))$$

is approximated for cell i by $\alpha D_0 \sum_j (m_j^{(n)} - m_i^{(n)})$ where j ranges over the neighbouring cells and α is a numerical constant.

Due to the discreteness of the numerical model, the long-time limits of Dutch and Belgian coins are 5% and 4%, respectively. Turning to the Figure 2 below, let us first consider the percentage of Dutch coins in the Netherlands (dotted line, right scale) in the diffusion-only scenario. As may be expected, the fraction of the local coins decrease monotonically over time towards its long-time equilibrium; this is also true for the airport-assisted cases (with faster decay rate) not shown here.

Next, we consider the percentages of Belgian coins in the Netherlands for three values of ε/D_0: 0 (solid line, left scale), 3×10^{-3} (long dashed line) and 0.1 (short dashed line). In all of these cases we see that the fraction of Belgian coins first increases before decaying to its long-time limit, with the "overshoot" being greater in the absence of long-distance travel. These are both to be expected, intuitively, due to the proximity of Belgium to the Netherlands.

To actually determine which of these qualitative curves best describe the reality, as well as to determine the actual timescales, we need to turn to the observed data. This turns out to be quite problematic, as we mentioned in Section 3, and for the above plot we have made a *very tentative* estimate for the timescale (in years).

This leaves us with several possibilities: First, assuming that our data accurately reflects the actual distribution of coins, the source terms over

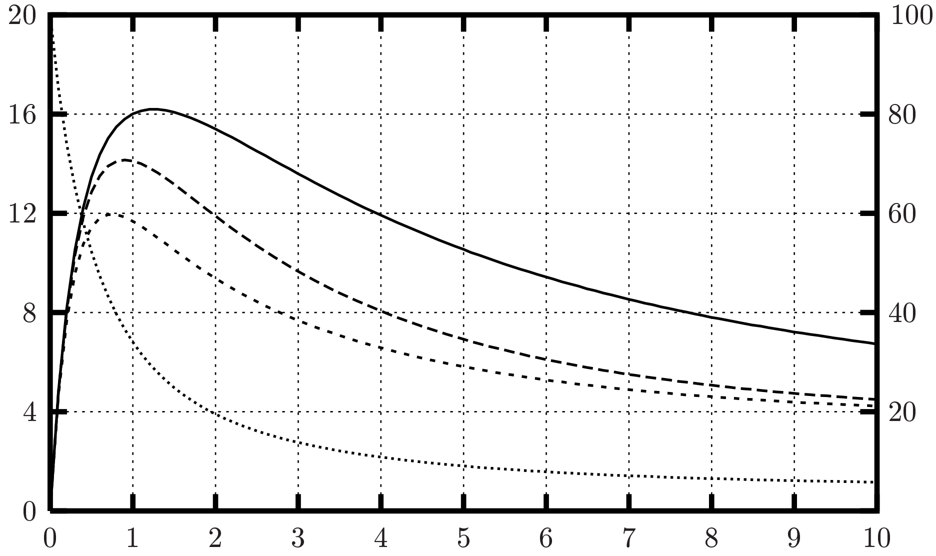

FIGURE 2. Percentages of Dutch and Belgian coins in the Netherlands (see text for details)

the timescale concerned may be important; this scenario can presumably be easily resolved by a simple check with the central banks. Second and still assuming the accuracy of our data, we may be forced to consider processes other than the mixing of coins, as no mixing model (regardless of the details) would produce an increase in the fraction of local coins; this possibility is difficult and, in our opinion, unappealing. Finally, the fluctuating tendencies may be caused by systematic bias in the data collection procedure; this possibility, too, is difficult to establish.

9. Conclusions*

In roughly a year half our wallet will be filled with foreign coins!

Acknowledgements

We thank the Euro Diffusion web site group for setting up the web site that is indispensable for the collection of data. We thank Jeanine Kippers from DNB (the Dutch national bank) for supplying additional information. We thank Richard Gill and Andreas Kyprianou for advice on the statistical part of the project the models respectively.