

Complaint

1 Deficiency in Supervision Received

According to the *Graduate Research Degrees Code of Practice 2010/11* (UCL Graduate School 2010), a student can expect their Principal Supervisor to 1) ‘establish a strict timetable of regular meetings with you at which all matters relating to your work can be discussed; such meetings should take place at least every two months’, 2) ‘provide you with prompt constructive criticism and reaction on submitted draft work to facilitate timely completion of your thesis’ and 3) ‘discuss with you the work accomplished and revise the objectives as the work proceeds in the light of progress made and external factors (such as newly published findings)’. All three regulations were breached. My Principal Supervisor, Philip Treleaven, Professor of Computing and Director of the Financial Computing Centre and previously a Pro-Provost, was a hands-off supervisor. Several years ago, when his primary research interest switched from finance to using 3D body scanning for clothes sizing, I was told to ask John Campbell, Professor of Computer Science, to be my new Principal Supervisor. I met with him, and he agreed to take on that role. His interests lie in multiagent systems, so to supplement the work on forecasting I added a chapter on multiagent modelling to my PhD thesis, which eventually led to criticisms such as ‘the scope of the thesis is far too broad’ (Tsang 2011). However, having only met him once, Professor Campbell retired. I found myself with no supervisor. Within the Department of Computer Science, Massimiliano Pontil and Mark Herbster were approached, and I met the former once, but neither were experienced or interested in the financial domain. There were no suitable candidates, so Professor Treleaven became my official supervisor again, but from then on I was effectively on my own.

My Subsidiary Supervisor was Dieter Girmes, in the Department of Statistical Science. According to UCL Graduate School (2010), a student can expect their Subsidiary Supervisor to ‘keep acquainted with the progress of your work, to be present at, at least, annual supervisory meetings and participate in the MPhil to PhD transfer process’, but I only met him once. Also, according to UCL Graduate School (2010), a student can expect their Subsidiary Supervisor to ‘provide for continuity of supervision in the event of the absence or departure of the Principal Supervisor’, but this never happened.

According to UCL Graduate School (2010), regarding their Departmental Graduate Tutor, a student can expect that ‘in the case of breakdown of rela-

tions between the Supervisor and the student and in other circumstances where a change of Supervisor is desirable, he or she assists in resolving any difficulties, and where necessary arranges the appointment of a replacement Supervisor'. The Departmental Graduate Tutor at the time, Angela Sasse, Professor of Human-Centred Technology, at first tried to help by meeting with me, but it was left to me to find a new supervisor. She later thought it easier to remove me from the department, so had my email account cancelled. I complained to the then Head of Department, Anthony Finkelstein, Professor of Software Systems Engineering and now Dean of the Faculty of Engineering Sciences. He resurrected my email, read my draft thesis, considered the situation, allowed me to stay and put me in touch with John Shawe-Taylor, Professor of Computational Statistics and Machine Learning, Director, Centre for Computational Statistics and Machine Learning and now Head of Department.

Professor Shawe-Taylor gave me three hours of unofficial but technical supervision (the only three hours I've ever had), was impressed with how quickly I did the work, used my thesis as an example to other students, told me to submit my thesis and gave me a good reference. He is one of the top people in the world in the area of machine learning, and thinks highly of me (see Exhibit A).

The *Academic Regulations and Guidelines for Research Degree Students: 2009-2010 Academic Session* (UCL Graduate School 2009a) states that 'a student whose thesis is referred by his or her examiners for substantial revision within 18 months ... will be entitled to have one draft of his or her thesis read and commented on by his or her principal supervisor and/or secondary supervisor before re-submission.' At this stage my principal supervisor, Professor Treleaven, only read the abstract and first chapter of my thesis.

2 Results of Examinations

My internal examiner was David Barber, Reader in Information Processing in the Department of Computer Science at UCL. Dr Barber did not originally wish to be an examiner as he was not interested in the financial domain, but eventually agreed. My external examiner was Edward Tsang, a professor in the School of Computer Science and Electronic Engineering at the University of Essex.

My first viva was conducted on 9 June 2009, and I was asked to resubmit my PhD thesis. On this basis the examiners also opted to conduct another viva, which was held on 4 July 2011.

2.1 The examiners could not reasonably have been aware of special circumstances which significantly affected the candidate's performance in the examination

The examiners would not have been aware that, despite paying full-time fees, I'd had only three hours of effective supervision, and at times, no supervisor, no

desk and no email, all in breach of regulations. This created stress, underlying anger and a sense of injustice.

2.2 The examination was not conducted in accordance with the relevant regulations/procedures

According to the document ‘Guidance on oral examinations for research degrees’ (UCL 2011) ‘in the interests of the candidate, UCL expects that the examination will be completed and examiners will have submitted the final joint report within three months of the despatch of the thesis to them.’ The time from the first submitted thesis being sent to the examiners (25 March 2009) to the day in which the examiners’ joint report was received (7 August 2011) was 4 months 13 days. The time from the resubmitted thesis being sent to the examiners (9 February 2011) to the date of the examiners’ joint report (4 July 2011) was 4 months 25 days. The examiners clearly did not have the candidate’s interests at heart, as both thesis submissions exceeded three months.

The regulations dictate that examiners are requested to each complete an independent preliminary report, which should be written by each examiner after reading the thesis but before conferring with the co-examiner about it, and all examiners’ reports should be sent to UCL within two working weeks of the oral examination (UCL Graduate School 2009b; UCL 2011). For the first viva, Professor Tsang failed to submit a preliminary report (see Exhibit B). For the second viva, both examiners submitted preliminary reports (Barber 2011; Tsang 2011), but they do not appear to be independent. In Dr Barber’s section on ‘General remarks’ under the first bullet point there are three examples of quotations from my thesis that are objected to. Professor Tsang’s section on ‘Weaknesses’ contains four quotations, three of which are the same as the examples given by Dr Barber. Dr Barber’s preliminary report is dated 1 July 2011 and two of his quotes were slightly wrong. Professor Tsang’s preliminary report is undated and his quotes are all accurate.

UCL (2011) states that ‘in the interests of the candidate, UCL expects that the oral examination will be conducted in a professional and open manner’ and ‘although there is no formal limit, typically an oral examination should last for around 2-3 hours and be a positive experience for the candidate, regardless of the examination outcome, as the examiners explore the original ideas and contribution to scholarship of the thesis.’ There was no preliminary small talk to relax the candidate, in fact there were no niceties whatsoever. From Dr Barber in particular there were no acknowledgements of the positive aspects of my thesis, just intellectual arrogance. The first viva was unpleasant, and the second even worse. The atmosphere was hostile throughout both vivas, the first viva was stuffy, and the second viva extremely hot (on one occasion my supervisor fetched bottles of water). The atmosphere was the most hostile and uncomfortable I have ever experienced. Despite never having fainted in my life, whilst standing at the white board I assumed that I was going to faint. Whilst being asked to solve an equation, all I could think of was what would happen if I fainted. Both vivas, the second of which was optionally conducted at the

discretion of the examiners, lasted four hours, making eight hours in total.

UCL (2011) states that in cases when examiners refer a thesis for re-submission for examination ‘the examiners’ joint report should include detailed guidance for the candidate and what is expected for them to achieve the standard for an award.’ I was not given *detailed* guidance of what was required to achieve the standard for an award. For example, the first point listed in Barber and Tsang (2009a), ‘The technical aspects in chapters 4 and 5 lack academic rigour and many of the experiments are lacking in sufficient detail to be reproduced’, is important yet far too vague to be useful.

2.3 Evidence that both examiners were biased or prejudiced against the candidate

Prior to the first viva, before any prejudice had time to take root, Barber (2009) starts with ‘My preliminary reading was that the thesis appears reasonably solid and well written’ and concludes with ‘in summary, I think the thesis is OK and don’t foresee too many difficulties.’ Unfortunately, even at this stage he could not resist countering positive aspects of the thesis with an unwarranted criticism of the candidate. For example, after noting that ‘there is a fair amount of material covering basic time-series methods’, he claims that ‘it would appear he is still somewhat an amateur in basic time-series methods’. I felt that there was an underlying theme through the entire process of both examinations of conceding that the thesis was acceptable, but taking a dislike to the candidate.

Barber (2009) continues ‘it’s not clear to me that the candidate could actually construct novel methods and techniques’. The artificial stock market and the DC-Fisher hybrid algorithm are both novel algorithms.

Dr Barber attempted to trivialise everything that I did. Barber (2009) describes the algorithms employed as ‘very basic off-the-shelf techniques’. Yet Chapter 4 of Sewell (2011i) contains one novel algorithm and one novel application of an algorithm. In Chapter 5 one algorithm is genuinely novel (DC-Fisher hybrid algorithm), one (the Fisher kernel) I derived half of the maths (calculation of the Fisher scores for the transmission probabilities) and then programmed the entire algorithm from scratch, two further algorithms were based on undocumented MATLAB code (Bayes point machine and DC algorithm) and perhaps just one algorithm could be described as ‘off-the-shelf’, albeit of an academic standard (LIBSVM).

Barber (2009) continues ‘how much he understands, for example, of the SVM technique is also questionable.’ Figures 5.1 and 5.2 in my thesis (Sewell 2011i) show two large screenshots of the only dedicated SVM software available for Windows which is freely available to download, and written by me (Sewell 2005a,b), which was ignored by the examiners.

Dr Barber, although clearly very bright and well qualified in the area of machine learning, has not published in the financial domain, so may be unaware of what is surprising, what is interesting and what is a contribution. During my first PhD viva Dr Barber admitted that he had not read all of my thesis. During this viva, he seemed curiously interested in intelligence and who was

the brightest. Dr Barber, despite being a Bayesian himself, seemed to dislike my Bayesian view of science (Sewell 2009a), I felt that we were philosophically opposed.

Professor Tsang was visibly angry that I had not cited his own work. Firstly I was shocked at how unprofessional this was, and secondly I considered it unfair as I didn't know that he would be my examiner when I submitted. He attached a list of references of his own publications to Barber and Tsang (2009a), which included a PhD thesis by one of his ex-students.

In all, I felt that the first viva focussed too much on the attitude of the candidate, rather than on his work.

Post-viva, Barber and Tsang (2009a) reported that 'the performance of the candidate during the examination was felt at times to be unscientific'. There were no incidences whatsoever of the candidate being 'unscientific' and the candidate is well acquainted with what science is (Sewell 2009a).

Upon reading my resubmitted thesis, and in reference to my work on an artificial stock market, in his preliminary report Barber (2011) states 'surely this is just a very specific model that the candidate has adopted, but cannot be claimed to represent all forms of technical analysis.' Firstly, this was explained, 'the technical analysts follow the technician's number one rule: they follow the trend, so the model fails to replicate some of the more complex procedures that chartists follow' (Sewell 2011i, p. 69). Secondly, I am well qualified to make that judgement, having passed the Society of Technical Analysts (STA) Diploma, and having been a member of the STA for over ten years.

Barber (2011) writes 'for example, there is an axiomatic description of probability on page 53 that begins with sigma-algebra and ends with stochastic processes. However, the candidate never makes much use of this formalism elsewhere in the thesis and one wonders why such incongruously formal statements are there. Is it just to impress the reader?' My work on stochastic processes derives from a lecture I gave at Royal Holloway university about five years ago (Sewell 2006) and the stochastic processes defined are used in Table 3.1 of Sewell (2011i, p. 56), which is necessary to characterize financial time series, which is the purpose of the chapter.

Barber (2011) asserts 'another slightly irritating example that I fear is for the impressionable reader is on page 81: "...outside quantum mechanics, no real system is truly linear." Are we to assume the candidate is speaking with some authority on this?' It was not an attempt at being pretentious, I mentioned the ubiquity of nonlinearity because it is a known fact (Meiss 2003), and to motivate using nonlinear methods.

Barber (2011) resented the sentence '...there is about a 39 per cent chance that stock market returns exhibit long memory ...' Firstly, I wrote no such thing. I wrote 'To conclude, there is about a 30 per cent chance that stock market returns exhibit long memory...' (Sewell 2011i, p. 27). Secondly, using all the information that one has available and putting a probability on a hypothesis is necessary if one wishes to conduct science (Sewell 2009a).

Barber (2011) wrote, 'The candidate was asked to explain the runs-test approach and, more generally, how to conduct a statistical test. The response

was unsatisfactory, with the candidate giving an incorrect description of the approach, relating the decision only in terms of the expected mean of the test-statistic under the null hypothesis. Whatever one feels about the general philosophy of statistical hypothesis testing, it would be expected that a scientist be familiar with the basic principles of the approach that he has taken.’ I programmed both of the runs tests from scratch, made the code available online (Sewell 2010b), and linked to it in my thesis (Sewell 2011i, p. 62). The second runs test could not previously be found online, so the coding and contribution is significant. I wrote hypothesis testing software and put it online seven years ago (Sewell 2004). I also wrote ‘Statistical inference (and what is wrong with classical statistics)’ (Sewell 2008c) three years ago (cited by Yang and Forbes (2011)) which evidences my comprehensive understanding of significance testing. I also spoke about significance testing at the Automated Trading 2008 conference in London (Sewell 2008a). During the viva the suggestion that I didn’t understand classical hypothesis testing riled me.

Barber (2011) complains that ‘whilst the thesis aims to be informal, at times there is a sudden switch to mathematical formality’. The thesis does not aim to be either formal or informal, but to communicate and argue the science in the best way possible: formal as and when necessary.

In his preliminary report, Tsang (2011) asserts that ‘in Chapter 3, results were presented. But no conclusions were drawn.’ However, conclusions are drawn in subsection ‘3.4 Conclusion and Summary’ (Sewell 2011i, p. 65).

Tsang (2011) continues ‘for example, for the runs test, it is not clear whether the candidate believed that the results support or refute the efficient market hypothesis.’ Yet Sewell (2011i, p. 65) explains that ‘the runs test uncovered highly significant patterns in DJIA daily returns that are inconsistent with a random walk, but not necessarily inconsistent with an efficient market.’

Tsang (2011) asserts that ‘Chapter 5 started with no free lunch theorem. It is not clear to me what impact it has to the rest of the chapter.’ Sewell (2011i, pp. 95–96) explains that ‘the no free lunch theorem for supervised machine learning states that in a noise-free scenario where the loss function is the misclassification rate, if one is interested in off-training-set error, then all algorithms are equivalent, on average. How well an algorithm performs is determined by how ‘aligned’ it is with the target function. The kernel constitutes prior knowledge that is available about a task, so the choice of kernel function is crucial for the success of all kernel algorithms.’

Tsang (2011) states that ‘Neely’s data set was used. This data set should be described.’ The data is described under the subheading ‘Data’ (Sewell 2011i, p. 96).

Barber (2011) objected to my claim that ‘one of the most exhaustive characterizations of financial time series available is undertaken’ (Sewell 2011i, p. 19). Similarly, Tsang (2011) objected to ‘an exhaustive characterization of financial markets was conducted’ (Sewell 2011i, p. 52). My literature review on the characterization of financial time series (Sewell 2011b) references 320 articles and had to be removed from the thesis and submitted as a Research Note as it took up too much space. Either a more thorough characterization of financial time

series should have been cited by the examiners (proof by contradiction), or the criticism withdrawn.

Barber (2011) and Tsang (2011) both took exception to my statement ‘a lot of the work presented in (most often IEEE) conference proceedings is of dubious quality’. Yet the claim is justified, in my literature review (Sewell 2011i, pp. 42 & 45) I explain why the methodology in the following two IEEE conference publications is flawed. Trafalis and Ince (2000) used supervised learning for regression to predict daily stock prices and, oddly, they forwent a validation set, and with the SVMs, set ϵ to zero, fixed C and repeated the experiment for various fixed settings of the kernel parameter, σ , giving rise to several results. Chen and Ho (2005) used an SVM for regression for forecasting the Taiwan Stock Exchange Capitalization Weighted Stock Index and, oddly, they considered price, rather than returns. Professor Tsang is a member of the IEEE, participates in its Computational Finance and Economics Technical Committee and is an Associate Editor of *IEEE Transactions on Evolutionary Computation*. Further, three of the four of Professor Tsang’s publications cited in my thesis are IEEE publications (Tsang and Martinez-Jaramillo 2004; Martinez-Jaramillo and Tsang 2009; Tsang and Isasi 2009). It is clear why Professor Tsang was unhappy with my criticism.

Tsang (2011) objected to my claim that ‘computer science is probably the fastest-changing discipline’ which is one of the statements that he considers ‘show immaturity of the candidate’. Further, both examiners made a big deal of this claim in my second viva, where I even explained why I thought it was a valid statement (computer science is relatively new, and Moore’s Law). According to David Hughes, Associate Professor and Department Chair, Computer Science Department, Brock University, who has been teaching the Introduction to Computer Science course at Brock since 1978, ‘Computer science is probably the most quickly changing of all subjects’ (Hughes 2002, p. 2).

An equation relating to technical analysis incorporating recency, with the most natural assignment of coefficients being that they sum to one and form a geometric sequence, involved solving the equation $c^3 + c^2 + c = 1$ (Sewell 2011i, p. 71). The equation has one real root, which is positive (0.544). Rather than considering the logic and the science (the need for a geometric sequence and summing to one), Dr Barber scented blood, so started challenging me to solve cubic equations, the maths of which has nothing to do with my thesis.

In their joint report, Tsang and Barber (2011) claim that “even basic points like the candidate’s claim that some of his algorithms ‘beat the market’ were very difficult for the candidate to explain.” This is a case of the examiner’s projecting their own ignorance onto the candidate. I have worked for a foreign exchange CTA plus three hedge funds, and am well aware of what ‘beat the market’ means, whilst Dr Barber is relatively new to the financial domain. In a foreign exchange market ‘beating the market’ simply means earning a positive return, but it is not a useful benchmark as it does not incorporate risk.

Immediately before my first viva I was reminded by my supervisor, Professor Treleaven, to defend myself. This made sense, the US word for a ‘viva’ being ‘defense’. However, to my surprise and horror I was told by Dr Barber to stop

being so defensive. There seemed to be too much emphasis on attacking my character, when we should have been discussing my thesis. Confusingly, before the second viva, both preliminary reports essentially said that I should pass, but only if I defend myself. “Provided that the above points are met and the candidate is able to defend the thesis adequately, I feel this could scrape through as a ‘low pass’” (Barber 2011) “Award of a PhD degree depends on satisfactory defence in the viva: the candidate showed weak scholarship in the previous viva” (Tsang 2011). Barber (2011) stated that ‘one of our major issues in the first viva was also the difficulty of verbally communicating with the candidate’ and ‘the candidate was particularly evasive in the first viva and he appeared to me at least to be more interested in ‘pseudo-science’ and philosophy, rather than in explaining and justifying the work he’d actually carried out’. Then during the second viva, I was told repeatedly and explicitly by Dr Barber that I was not allowed to defend myself. After the second viva, Tsang and Barber (2011) wrote ‘in general, his manner was very evasive in the viva’. If I appeared evasive, it was because I was absolutely livid at being told that I couldn’t defend myself. I was practically speechless with anger.

Tsang and Barber (2011) asserted that ‘we have not been able to identify a single topic that has been sufficiently well described to be replicated independently.’ The detail provided in the step-by-step procedure I wrote for the Fisher kernel methodology (Sewell 2011i, pp. 94–95 & pp. 97–98) far exceeds that expected in a journal publication. In most cases I also put the software online, so that my experiments can be replicated by anyone.

Two of the criteria that the examiners deemed satisfactory on the first summary report (Barber and Tsang 2009b), namely ‘is an integrated whole and presents a coherent argument’ and ‘demonstrates research skills relevant to the thesis’ were left unchecked on the second joint report (Tsang and Barber 2011). Either spending six months addressing all of the examiners’ issues listed in their first joint report (Barber and Tsang 2009a) has made the thesis worse in their own eyes, or over time they came to dislike the candidate, i.e. prejudice had time to set in.

Tsang and Barber (2011) wrote, ‘Despite the large number of tangential and informal reports written by the candidate, we managed to identify (with the candidate’s help) only three works that appeared to be peer reviewed in any way, none of these having any significant connection with the material presented in the thesis.’ This is a complete misrepresentation of the facts. There were no ‘informal reports’ written by the candidate referenced in Sewell (2011i), but six official UCL Department of Computer Science Research Notes (Sewell 2011b,c,d,e,f,g) which were written at the explicit request of the examiners (Barber and Tsang 2009a) (and ignored or dismissed in the second viva), and no less than seven peer reviewed articles (Sewell 2008b; Sewell and Yan 2008; Yan, Sewell and Clack 2008; Sewell 2009b, 2010a, 2011h,a). Of these, Sewell (2011h) describes the evolution of the heuristics and biases used in the artificial stock market described in Chapter 4 (this work was in the thesis before I was asked to remove it by my supervisor), Sewell (2009b) relates directly to the experiment on performance measurement in Chapter 4, and Yan, Sewell and Clack (2008)

provides a head-to-head evaluation of GP and SVM forecasting, mirroring the work in Chapter 5, and drawing the same conclusion.

Tsang and Barber (2011) wrote ‘the candidate has some difficulty with this basic algebraic manipulation. It’s disturbing that the candidate claims to understand sigma-algebra and Levy processes, yet cannot easily rearrange a simple equation.’ Firstly, this is untrue. I was asked to rearrange an equation on the white board, did so correctly, and Dr Barber acknowledged that it was correct. Secondly, it is insulting. I did maths as an undergrad (BSc (Hons) Mathematics from the University of Bristol) and resent being tested on school-level maths. Thirdly, it is unnecessary. I couldn’t possibly have written the model in the first place if I was not au fait with exponents and logarithms. It is also facetious, and appeared to be motivated by little more than an attempt to humiliate the candidate.

Tsang and Barber (2011) are of the opinion that ‘the material in chapters 3, 4, and 5 could potentially be published, although it would require some additional work to improve clarity and inferences drawn from the work’, whilst UCL Graduate School (2009b) requires that a PhD thesis ‘contains an element which might, after any necessary revision, merit publication in a medium appropriate to the discipline (for example as a monograph or as a number of articles in learned journals’. On that basis, the thesis is consistent with what is required of a PhD. Besides, the PhD thesis from the Department of Computer Science at UCL that I was given as a good relevant benchmark, titled *The Application of Intelligent Systems to Finance and Business* (Viner 1998), contains just one self-citation, a conference paper (Viner and Goonatilake 1996).

After the first viva I offered to submit two chapters as papers to conferences, but Professor Tsang insisted that I submit to a journal. The problem was that most journals could not provide feedback in the timeframe available.

The examiners ignored all of the work that they’d requested I do in Barber and Tsang (2009a), work that I had just taken six months off work to complete. I was equally careful not to change work that they were apparently satisfied with.

The examiners also completely ignored numerous genuine contributions. For example, I show that if one conditions on price, then technical analysis is self-fulfilling, and if one conditions on time, then technical analysis is self-destructive (Sewell 2011i, p. 31). I also show that the foreign exchange markets are the most predictable, futures markets intermediate and stock markets the least predictable (Sewell 2011i, pp. 105 & 109), cited by McBride (2010). I explain that technical analysis is representativeness (Sewell 2011i, pp. 30 & 70). I give the necessary conditions where one can use K-fold cross-validation on a time series (Sewell 2011i, p. 83).

In academia computer scientists often dismiss programming, because they claim that it is the algorithm that is of interest. However, programming an algorithm from scratch and putting it in the public domain for scrutiny evidences that the candidate understands the algorithm explicitly. Further, there is no reason why computer science, an engineering science within the Faculty of Engineering Sciences, should not be judged from an engineering perspective. I wrote

no less than ten programs over the course of my thesis, all with source code in the public domain: Rescaled Range Analysis in C++ and Visual Basic for Excel, Runs Test in Visual Basic for Excel, Performance Measurement Calculator in PHP and Visual Basic for Excel, Performance Metric Analysis in Visual Basic for Excel, SVMdark in C for Win32, winSVM in C++ for Win32, Fisher Kernel in C++, Monte Carlo Portfolio Optimization in Visual Basic for Excel, Kelly Criterion in PHP and Visual Basic for Excel, and Order Book Reconstruction in C#.

After the second viva, given that most of the recommended changes that the examiners made involved deleting sentences that they didn't approve of, or improving explanations of the methodology generally, they could have offered minor corrections in three months.

Professor Tsang is the Director of the Centre for Computational Finance and Economic Agents (CCFEA), University of Essex, and the Deputy Director is Dr Wing Lon Ng. Dr Ng has already somehow cited my thesis in a journal article he wrote with one of his PhD students, Abdalla Kablan (Kablan and Ng 2011). Mr (now Dr) Kablan won the David Norman Prize at CCFEA for this paper.

References

- BARBER, David, 2009. Examiner's preliminary report for a research degree level candidate. 8 June 2009.
- BARBER, David, 2011. Examiner's preliminary report for a research degree level candidate. 1 July 2011.
- BARBER, David, and Edward TSANG, 2009a. Examiners' joint report for a research degree level candidate. 15 June 2009.
- BARBER, D., and E. TSANG, 2009b. Examiners' summary report for a PhD candidate. 9 June 2009.
- CHEN, Kuan-Yu, and Chia-Hui HO, 2005. An improved support vector regression modeling for Taiwan Stock Exchange market weighted index forecasting. *In: Mingsheng ZHAO and Zhongzhi SHI, eds. ICNN&B '05: International Conference on Neural Networks and Brain, 2005, Volume 3.* IEEE, pp. 1633–1638.
- HUGHES, David, 2002. *Fundamentals of Computer Science Using Java*. Sudbury, MA: Jones and Bartlett.
- KABLAN, Abdalla, and Wing Lon NG, 2011. Intraday high-frequency FX trading with adaptive neuro-fuzzy inference systems. *International Journal of Financial Markets and Derivatives*, **2**(1/2), 68–87.
- MARTINEZ-JARAMILLO, Serafin, and Edward P. K. TSANG, 2009. An heterogeneous, endogenous and coevolutionary GP-based financial market. *IEEE Transactions on Evolutionary Computation*, **13**(1), 33–55.
- MCBRIDE, Sarah, 2010. *Markets and conspiracy theories*. Washington, DC: NPR. 16 August 2010.
- MEISS, Jim, 2003. Sci.nonlinear FAQ 2.0. <http://amath.colorado.edu/faculty/jdm/faq-Contents.html>.
- SEWELL, Martin, 2004. Significance test. <http://www.stats.org.uk/significance-test.php>.
- SEWELL, Martin, 2005a. SVMdark. <http://www.cs.ucl.ac.uk/staff/M.Sewell/svmdark/>.
- SEWELL, Martin, 2005b. winSVM. <http://www.cs.ucl.ac.uk/staff/M.Sewell/winsvm/>.
- SEWELL, Martin, 2006. Stochastic processes. <http://www.cs.ucl.ac.uk/staff/M.Sewell/stochastic-processes.pdf>. Lecture given to Department of Economics, Royal Holloway, University of London, Egham.

- SEWELL, Martin, 2008a. Optimization and methods to avoid overfitting. Talk presented at the Automated Trading 2008 conference, London, 15 October 2008.
- SEWELL, Martin, 2008b. Gender (*sic*) equality (*sic*). *Opticon1826*, 4.
- SEWELL, Martin, 2008c. Statistical inference (and what is wrong with classical statistics). <http://www.stats.org.uk/statistical-inference/>.
- SEWELL, Martin, 2009a. The demarcation of science. <http://science.martinsewell.com/science.pdf>.
- SEWELL, Martin, 2009b. Decision making under risk: A prescriptive approach. In: *Proceedings of the 2009 Symposium; Academy of Behavioral Finance and Economics, September 23-25, 2009, Chicago, Illinois*. Montrose, CA: Academy of Behavioral Finance and Economics, pp. 9.
- SEWELL, Martin, 2010a. Emotions help solve the prisoner's dilemma. Behavioural Finance Working Group Conference: Fairness, Trust and Emotions in Finance, London, 1-2 July 2010.
- SEWELL, Martin, 2010b. Runs test. <http://www.stats.org.uk/runs-test.xls>.
- SEWELL, Martin, 2011a. Policy via science. Discussion Paper 2011-01, E3 Foundation, Cambridge.
- SEWELL, Martin, 2011b. Characterization of financial time series. Research Note RN/11/01, University College London, London.
- SEWELL, Martin, 2011c. Ensemble learning. Research Note RN/11/02, University College London, London.
- SEWELL, Martin, 2011d. Fund performance. Research Note RN/11/03, University College London, London.
- SEWELL, Martin, 2011e. History of the efficient market hypothesis. Research Note RN/11/04, University College London, London.
- SEWELL, Martin, 2011f. Money management. Research Note RN/11/05, University College London, London.
- SEWELL, Martin, 2011g. The Fisher kernel: A brief review. Research Note RN/11/06, University College London, London.
- SEWELL, Martin, 2011h. The evolution of entrepreneurs and venture capitalists. In: Rassoul YAZDIPOUR, ed. *Advances in Entrepreneurial Finance: With Applications from Behavioral Finance and Economics*. New York: Springer, Chapter 11, pp. 205-217.

- SEWELL, Martin Victor, 2011i. *The Application of Intelligent Systems to Financial Time Series Analysis*. Ph. D. thesis, University College London, London.
- SEWELL, Martin V., and Wei YAN, 2008. Ultra high frequency financial data. In: Maarten KEIJZER, ed. *Proceedings of the 2008 GECCO Conference Companion on Genetic and Evolutionary Computation*. New York: ACM, pp. 1847–1850.
- TRAFALIS, Theodore B., and Huseyin INCE, 2000. Support vector machine for regression and applications to financial forecasting. In: Shun-Ichi AMARI, et al., eds. *IJCNN 2000, Proceedings of the IEEE-INNS-ENNS International Joint Conference on Neural Networks. Volume 6*. Los Alamitos: IEEE Computer Society, pp. 6348.
- TSANG, Edward, and Pedro ISASI, 2009. Editorial special issue: Computational finance and economics. *IEEE Transactions on Evolutionary Computation*, **13**(1), 1–2.
- TSANG, Edward P.K., and Serafin MARTINEZ-JARAMILLO, 2004. Computational finance. *IEEE Computational Intelligence Society Newsletter*, , 8–13.
- TSANG, Edward P. K., 2011. Examiner’s preliminary report for a research degree level candidate.
- TSANG, Edward P. K., and David BARBER, 2011. Examiners’ joint report for a PhD candidate. 4 July 2011.
- UCL, 2011. *Guidance on oral examinations for research degrees*.
- UCL Graduate School, 2009a. *Academic Regulations and Guidelines for Research Degree Students: 2009-2010 Academic Session*. Section 1: Programme of Study.
- UCL Graduate School, 2009b. *Academic Regulations and Guidelines for Research Degree Students: 2009-2010 Academic Session*. Section 2: Examinations.
- UCL Graduate School, 2010. *Graduate Research Degrees Code of Practice 2010/11*.
- VINER, James, 1998. *The Application of Intelligent Systems to Finance and Business*. Ph. D. thesis, University College London, London.
- VINER, James, and Suran GOONATILATE, 1996. A study of a genetic algorithm based continually adaptive trading engine. In: *Proceedings of the International Conference on Intelligent Technologies in Human-Related Sciences*. AMSE.

- YAN, Wei, Martin V. SEWELL, and Christopher D. CLACK, 2008. Learning to optimize profits beats predicting returns — comparing techniques for financial portfolio optimisation. *In: Maarten KEIJZER, ed. Proceedings of the 10th Annual Conference on Genetic and Evolutionary Computation.* New York: ACM, pp. 1681–1688.
- YANG, Xin-She, and Alistair B. FORBES, 2011. Model and feature selection in metrology data approximation. *In: Emmanuil H. GEORGOULIS, Armin ISKE, and Jeremy LEVESLEY, eds. Approximation Algorithms for Complex Systems, Volume 3 of Springer Proceedings in Mathematics.* Springer, pp. 293–307.