UNIVERSITY OF CAMBRIDGE

Electricity Policy Research Group

Fuel Panics

insights from spatial agent-based simulation

EPRG Working Paper 1305

Cambridge Working Paper in Economics 1309

Eben Upton and William J. Nuttall

Abstract

The United Kingdom has twice suffered major disruption as a result of fuel panics first in September 2000 coincident with a wave of fuel protests and more recently in March 2012 following politcal warnings of possible future supply chain disruption. In each case the disruption and economic consequences were serious. Fuel distribution is an example of a supply chain. Approaches to supply-chain planning based on linear programming are poorly suited to modelling non-equilibrium effects, while coarse-grained system dynamics models often fail to capture local phenomena which contribute to the evolution of global demand. In this Paper, we demonstrate that agent-based techniques offer a powerful framework for co-simulation of supply chains and consumers under conditions of transient demand. In the case of fuel panic crisis, we show that even a highly abstract model can reproduce a range of transient phenomena seen in the real world, and present a set of practical recommendations for policymakers faced with panic-buying.

Keywords Fuel Panics, Agent Based Simulation, Supply Chain

JEL Classification C15, C63, H30, J48, L91, R40

Contact william.nuttall@open.ac.uk
Publication April 2013
Financial Support N/A

Fuel Panics
insights from spatial agent-based simulation

Eben Upton and William J. Nuttall*

Judge Business School, University of Cambridge
Cambridge, CB2 1AG, UK

* Present address: Department of Design, Development, Environment and Materials,
The Open University, Walton Hall, Milton Keynes, MK7 6AA, UK

Abstract

The United Kingdom has twice suffered major disruption as a result of fuel panics first in September 2000 coincident with a wave of fuel protests and more recently in March 2012 following politcal warnings of possible future supply chain disruption. In each case the disruption and economic consequences were serious. Fuel distribution is an example of a supply chain. Approaches to supply-chain planning based on linear programming are poorly suited to modelling non-equilibrium effects, while coarse-grained system dynamics models often fail to capture local phenomena which contribute to the evolution of global demand. In this Paper, we demonstrate that agent-based techniques offer a powerful framework for co-simulation of supply chains and consumers under conditions of transient demand. In the case of fuel panic crisis, we show that even a highly abstract model can reproduce a range of transient phenomena seen in the real world, and present a set of practical recommendations for policymakers faced with panic-buying.

1 Introduction

In this paper, we argue that the methodology of *agent-based simulation* can provide a unified framework for modelling supply chains and the social and spatial dynamics which give rise to transient demand. We begin with a brief overview of the literature relating to supply-chain modelling, agent-based simulation and models of individual and group behaviour. Taking as our example the UK panic-buying of petrol as occured in September 2000 and March 2012, we show that even a highly abstract model can reproduce a range of transient phenomena seen in the real world; we use our model to develop insights into the causes of these phenomena, and to evaluate the effectiveness of several proposed responses, and offer a set of recommendations for policymakers faced with panic-buying. We conclude by characterising the varieties of problem to which we believe agent-based techniques are best suited.

2 Agent-based simulation

Agent-based simulation is a technique for modelling complex systems through the interaction of a population of autonomous software objects. Individual agents interact with their environment and each other, typically following a simple behavioural schema, which can be represented by a finite-state automaton. Models of physical systems often embed agents in a multi-dimensional space, with inter-agent connectivity wholly or partially determined by proximity; perhaps the most famous such model is Conway's Game of Life [1], in which cells in a two-dimensional grid live or die based on the local population density. The macro-scale behaviour of the system is an emergent property of the agents' schemata, and of their connectivity. This behaviour can be very complex indeed; the Game of Life is known to be Turing-powerful (formally equivalent to a universal computer), while Wolfram [2] contends that the fundamental laws of nature themselves can be understood in terms of the interaction of simple agents.

It should be emphasised that the primary value of agent-based simulation is in illuminating the dynamics of a real-world system in a controlled environment amenable to experimentation, rather than in predicting the system's exact future behaviour. Even in a purely deterministic system, imprecision in measurements of initial conditions and the modelling of agent behaviour generally causes the simulation to diverge from reality. Axelrod and Tesfatsion [3] cite three goals commonly pursued by researchers:

- **Empirical understanding.** How are macro-scale features of the system determined by the micro-scale behaviour of the participants?
- **Normative understanding.** Can we discover micro-scale interventions which will produce favourable macro-scale outcomes in the real system?
- **Methodological advancement.** How can we provide better tools to support simulation and the validation of the results against observational data?

Agent-based simulation has been applied to many problems in the social sciences. A seminal example is the Schelling Segregation Model [4], which describes how communities of reasonably tolerant individuals may spontaneously segregate along racial lines. Epstein [5] presents models of rebellion and communal violence, in which a small group of police attempts to control the behaviour of a much larger population of citizens. Axelrod [6] and Epstein [7] separately examine the evolution of norms in a society where citizens attempt to establish and conform to common patterns of behaviour; Epstein's approach is somewhat more sophisticated than Axelrod's purely game-theoretic treatment, introducing both a physical embedding of agents in a one-dimensional space with periodic boundary conditions, and a notion of "adaptive sampling", in which agents attempt to minimise the work done in detecting the norm. Hamilton et al [8] describe an agent-based model of technology diffusion, in which users adopt a new technology based on both their perception of its performance and the behaviour of their immediate neighbours; Zhang and Nuttall [9] take a similar approach, incorporating a model of intention (the "theory of planned behaviour", described in greater detail below) and a social network overlay with small-world properties.

Drogoul et al [10] present a model-based methodology for the design of agent-based simulations, summarised in Figure 1. They identify three distinct roles involved in the design process, at successively higher levels of remove from the target system, each with its own definitions of the terms model and agent. The *thematician*, or domain specialist, thinks in terms of a domain model, populated by real agents, which serve as metaphors for entities in the real environment; his description of agent behaviour encapsulates his micro-knowledge of entity behaviour, in the form of observations and assumptions. The *modeller* is responsible for formalising the domain model, to produce a design model, populated by conceptual agents; this model specifies interactions between agents in a notation with a formally defined operational semantics, and can be seen as an abstraction both of the domain model and of the subsequent implementation. Finally, the *computer scientist* defines and implements a corresponding operational model composed of computational agents. Feedback loops involve the modeller in internal validation of the operational model against the design model, and the thematician in analysis and external validation of simulation results against his macro-knowledge of the target system.

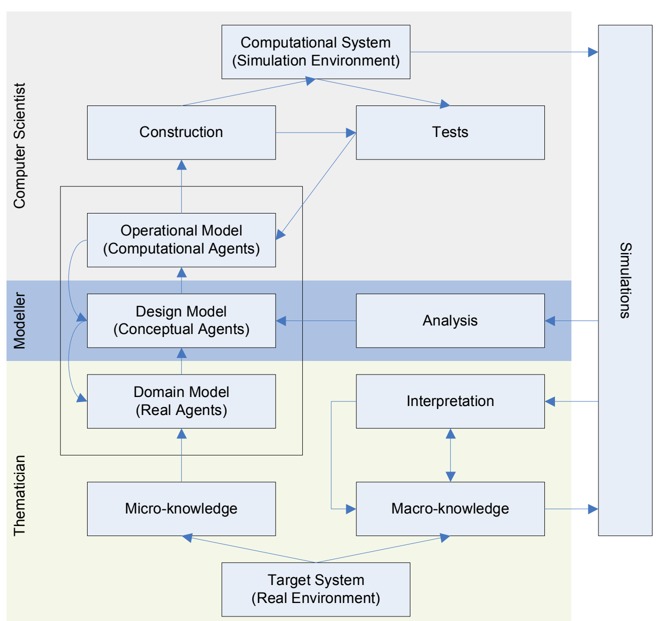

Figure 1: Methodology for the design of agent-based simulations (source: Drogoul et al)

2.1 Individual behaviour

The fidelity of a simulation which incorporates agents representing individuals will naturally depend on the precision with which it models the decision-making process. The fields of economics and behavioural science provide us with two complementary models of human decision-making.

Mainstream economic decision theory rests on the concept of rational choice, in which an individual's preferences for particular actions are captured by a scalar utility function, and the central goal of decision-making is the maximisation of expected subjective utility; an actor following this decision rule is said to be a rational maximiser. Preferences may be fixed or time-varying, and actors may possess total knowledge and unlimited computational ability or only partial knowledge and/or limited computational ability; the latter case corresponds to Simon's notion of bounded rationality [11]. Rational choice theories are traditionally formulated in terms of "games against nature"; game theory can be seen as an extension to the case in which one's opponent is also believed to be a rational actor.

The widespread tendency, which Lazear [12] characterises as "economic imperialism", to apply rational choice theories to problems in the social sciences outside of economics remains controversial. A significant body of empirical evidence exists to suggest that under many circumstances humans do not in fact behave as rational maximisers. Thaler [13] has shown, for example, that individuals expect to be paid more to give up an object than they would pay to acquire it; this behaviour is referred to as divestiture aversion, or the endowment effect. The field of behavioural economics attempts to construct improved models of human behaviour by incorporating insights drawn from experimental psychology, notably the existence of the endowment effect, and other cognitive biases including loss aversion and zero-risk bias. Kahneman and Tversky's [14] prospect theory attempts to account for these effects and to explain the associated deviations from strict rationality by:

- assigning value to gains and losses rather than to absolute assets; and
- applying non-linear weighting functions to outcome probabilities and payoffs.

Trepel et al [15] provide evidence that these components of prospect theory have a basis in the physiology of the human brain, while Thaler and Sunstein [16] discuss how governments may exploit biases to "nudge" citizens into adopting desired behaviours. From a modelling perspective, the theory is appealing because it retains the simplicity of rational choice theory while offering greater predictive accuracy.

In the field of behavioural science, Ajzen's theory of planned behaviour [17] aims to explain an individual's behaviour in specific contexts in terms of three interrelated factors:

- **Attitude.** Whether the individual is well-disposed towards engaging in the behaviour.
- **Subjective social norms.** The perceived degree of social pressure to engage in the behaviour.
- **Perceived behavioural control.** The anticipated ease of engaging in the behaviour.

As in the earlier theory of reasoned action, these factors serve as the antecedents of intention, a term which represents the degree of effort the individual is prepared to expend to engage in the behaviour. Whether the behaviour takes place then depends on both intention and perceived behavioural control. The latter contributes directly because, holding intention constant, an individual who believes that the intended behaviour is achievable is more likely to successfully engage in it, and because perceived control serves as a good proxy for actual control, which in turn is a prerequisite for action. The sensitivity of each source of intention to external factors is determined by endogenous parameters which correspond to the individual's personality traits or beliefs.

The theory of planned behaviour provides the modeller with a framework for incorporating additional non-economic factors into an agent's schema. These factors may prevent an agent from engaging in an otherwise economically favourable activity if it violates a perceived norm or is seen as being too difficult to accomplish, or vice versa. In Zhang and Nuttall's simulations of technology adoption, individual attitude is determined by economics, while subjective social norms and perceived behavioural control are influenced by input from the agent's social network and the actions of regulators. The favourability of each possible action is determined by a weighted sum of these inputs, and at each time step the agent takes the most favourable action.

2.2 Group behaviour

Models of group behaviour seek to bridge the gap between the comparatively simple behaviour of individuals and the often extremely complex dynamics of groups of interacting individuals. A relevant example is the theory of *information cascades*, put forward by Bikhchandani et al [18] to explain the high degree of localised conformity observed in human society. These cascades can lead rational individuals to mimic the behaviour of others rather than relying on their own private information. They establish a model in which a series of individuals makes a binary decision about a single hidden variable drawn from the set $\{YES, NO\}$, based on a "noisy" signal (a piece of possibly incorrect private information) and the decisions of all previous participants, and show that:

- the system collapses into a YES or NO cascade, with individuals ignoring their information;
- the system is liable to collapse into an incorrect cascade (YES if the answer is NO or vice versa);
- the less noisy the signal, the faster the collapse and the greater the likelihood of a correct result;
- cascades are fragile, and can be disrupted by the appearance of new public information.

The key realisation is that a rational individual will attempt to extract information about the true value of the hidden variable from the actions of previous participants; even a small imbalance between the observed YES and NO counts will overwhelm the individual's private information, triggering a self-reinforcing cascade. Because a cascade prevents private information from affecting observable actions, the n^{th} participant in a cascade has no more information from his predecessors than the first participant; as a result, a single piece of new public information is enough to tip the balance and destroy the cascade.

All models of group behaviour rely on implicit or explicit assumptions about the structure of the connection graph linking individuals. Empirical studies, such as that by de Solla Price [19], have shown that for many human social networks this graph has the scale-free property, in which the number of connections from a node (its *degree*) obeys a power-law distribution, and low-degree nodes are much more tightly clustered than high-degree ones. The small number of high-degree nodes act as hubs, dramatically reducing the average path length compared to a randomly connected graph; scale-free networks are therefore resilient to random disruption, as deleting a random node is unlikely to remove a hub, but are vulnerable to more targeted attacks. In a social network, clusters of low-degree nodes correspond to social groups (neighbours, co-workers), while high-degree nodes correspond to well-connected individuals, for example those in the media, who have the ability to directly influence individuals in many groups.

Several algorithms have been proposed for generating artificial scale-free networks for the purposes of simulation. We highlight two of these here:

- The **Barabsi-Albert (BA) model** [20] generates a network incrementally by a process of preferential attachment. Beginning with a pair of connected nodes, new nodes are added to the graph one by one. Each new node is connected to m existing nodes; the probability of selecting a node as the target for a connection is proportional to its current degree. The BA model produces graphs with a realistic degree distribution, but does not accurately replicate the clustering characteristics of real-world networks.
- The **Watts-Strogatz (WS) model** [21] begins with a regular ring lattice; this consists of a circle of nodes, each of which is connected to its n nearest clockwise and counter-clockwise neighbours. Each edge is then rewired with probability p to point to a randomly selected node, subject to the constraint that rewiring does not introduce loops or duplicate links. Varying p interpolates between the ring lattice and an Erdős-Rényi random graph [22]. The WS model produces graphs with realistic clustering characteristics, but with an unrealistically homogeneous degree distribution.

We will see in Section 3 how the WS model can be applied to simulations in which agents are embedded in a one-dimensional space with periodic boundary conditions.

3 Fuel Panics

We apply the methodology of agent-based simulation to the panic-buying of petrol. We begin with an overview of the UK petroleum retailing sector, and provide a brief history of the two major outbreaks of panic-buying in the UK, during the fuel protests of September 2000 and amidst mere fear of supply disruption in March 2012. We then develop a highly abstracted design model of fuel distribution and consumer behaviour, and conduct several experiments using the model to gain an insight into the dynamics of the real system. Finally, we present a set of recommendations for policymakers based on the outcome of our experiments.

3.1 UK petroleum retailing

Over the past decade, the UK petroleum retailing sector has experienced a process of consolidation, driven by strong competition and by the fixed per-site costs of compliance with environmental regulations. As of July 2010 there were approximately 9,000 filling stations in the UK [23]. Supermarkets account for 1,200 sites, or 14% of the total. Average annual throughput is 2.2 million litres per year for a standalone site, or 11 million litres per year for a supermarket site. Filling stations are resupplied by road tanker from one of around 50 major depots, which in turn are resupplied from refineries by pipeline, rail and sea. Tankers have a capacity of around 30,000 litres; assuming dedicated deliveries, standalone sites must be resupplied on average once every five days, and supermarket sites daily. Resupplying stations more frequently with partial tanker-loads can reduce inventory requirements somewhat; Boctor et al [24] describe several techniques for solving the resulting trip-packing problem.

Precise data on filling-station tank capacity are hard to obtain, as the majority of standalone filling stations are owner-operated small businesses. Table 1 shows data for four stations, taken from planning applications [25] and business-for-sale advertisements [26, 27, 28]; assuming that these are representative, we obtain an estimated total capacity across the filling-station network of 620 million litres.

Location	Type	Annual throughput (l)	Tank capacity (l)	Days' fuel
Suffolk	Supermarket	11,000,000 (est)	135,000	4
Surrey	Standalone	1,820,000	52,000	10
S Glamorgan	Standalone	1,200,000	62,000	19
Kent	Standalone	1,020,000	27,000	10

Table 1: Filling station throughput and tank capacity (source: various)

At the end of 2009, there were 34.3 million vehicles registered in the UK [29]. Assuming an average fuel-tank capacity of 55 litres, we obtain an estimated total capacity across all vehicles of 1.9 billion litres, or roughly three times the capacity of the filling-station network. Dividing this figure by the daily filling-station throughput of 83 million litres gives a mean time between refuelling of 23 days.

To date, modelling of petroleum retailing has concentrated on accounting for the observed spatial and temporal variation in prices. Anderson [30] presents a model in which competitive dynamics combine with consumers' tendency to refuel early if offered an attractive price to produce short-period price oscillations. Heppenstall et al [31] have developed a family of agent-based simulations which model the geographical distribution of petrol prices, taking into account retailer strategies and consumer preferences.

3.2 September 2000

Periodic direct-action protests over the cost of fuel have been a feature of political discourse throughout Western Europe over the past two decades. The first significant example in the UK took place in September 2000, when a coalition of farmers and road hauliers blockaded refineries and depots across the British mainland in protest against the high price of diesel fuel brought about by the fuel duty escalator, and competition from foreign hauliers. Robinson [32] provides a chronology of the protests, which may be summarized as follows:

7 September 150 protesters blockade Stanlow refinery in Cheshire for several hours, before being moved on by police. Only one of 60 scheduled Shell tankers departs.

8–9 September Protests spread to other depots. Some filling stations closed due to lack of fuel in Northern England. Isolated reports of panic-buying of fuel.

11 September Most refineries and depots effectively closed. Government secures emergency powers to protect fuel supplies for essential services.

12 September 25% of filling stations closed. Widespread panic-buying of fuel. Government announces deliveries will return to normal within 24 hours.

13 September 90% of filling stations closed. Deliveries at less than 4% of normal levels. NHS trusts begin to cancel non-essential operations. Royal Mail deliveries cease and schools close in many areas. Panic-buying spreads to groceries; Tesco and Sainsbury's introduce rationing.

14 September Government begins to deploy military tankers, and provides police escorts for commercial tankers. Protest leaders call off blockades in face of changing public opinion.

16 September Fuel deliveries at 135% of normal levels. Filling stations begin to reopen.

The economic impact of the protests was estimated at £1bn by the Institute of Directors [33]. Department for Transport statistics [34] indicate that motorway traffic had fallen to 39% below normal levels by 14 September, and did not fully recover until the end of the month; we see from Figure 2 that the recovery was not smooth, but instead exhibited marked oscillations. On a political level, the protests were moderately successful, forcing the government to abandon the fuel duty escalator and to introduce cuts in duty on ultra-low-sulphur petrol as part of the November 2000 pre-Budget report.

Of particular interest is the role of panic-buying in depleting fuel reserves. Based on the figures in Table 1, and assuming that car and station fuel tanks were at random levels at the start of the protests, we might expect it to have taken five days from the cessation of supply for 50% of standalone filling stations to close; in reality, 90% of filling stations were closed within two days. The high ratio of aggregate vehicle-tank to station-tank capacity allowed opportunistic early refuelling to overwhelm the system, creating a situation in which the remaining fuel was concentrated in the tanks of "lucky" drivers, while their unlucky counterparts were forced off the road. It is the opinion of the authors that attempts to discourage panic-buying were hindered by three key factors:

- the lack of a pre-existing organised system for enforcing rationing;

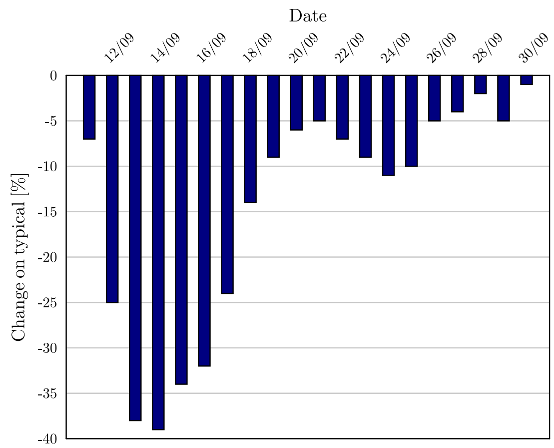

Figure 2: Percentage change on average motorway traffic (source: Department for Transport)

- retailers' reluctance to increase prices, for fear of being accused of gouging;
- the well-documented low price elasticity of demand for fuel [35].

3.3 March 2012

On 25 March 2012 British media, including the Guardian newspaper, reported that the Minister for the Cabinet Office, Francis Maude MP, had recently called on the tanker drivers' union Unite to seek an agreement in an industrial dispute in which a ballot for strike action was being planned [36]. All parties, including the minister, were acutely aware of the experiences of September 2000 and reports claimed that the Minister said in his remarks that the government had ”learned the lessons of the past” [36]. Given what was to follow, it was far from clear that they had. Over the following days public anxiety grew and demand rose at filling stations. After two days on 27 March the first filling stations started to run out, buy the following day it was a full sclae fuel panic crisis [37].

On Wednesday 28 March, Francis Maude appeared on Sky News and offered advice that was later to be criticised for being dangerous. He said: ”When it makes sense, a bit of extra fuel in a jerry can in the garage is a sensible precaution to take” [38]. Despite no strike action, nor any other disruption to the supply chain, the mere fear of disruption fanned by political statements had been sufficient to prompt a serious wave of panic buying across the whole country.

One key difference from September 2000 relates to the ubiquity of internet communications and hence opportunities to understand unfolding events. As a consequence of technological improvements over the intervening twelve years bottom up approaches to data gathering were now possible, which when combined with data visualisation tools allowed the panic to be assessed almost as it unfolded. An example using Google Fusion tables was developed by staff at the Guardian [37]. This example was developed for public consumption, in some sense closing the circle to those that provided the original data.

The fact that the the March 2012 crisis did not have its origins in a real supply chain problem is interesting, especially when considering the closing phase of the crisis. While reassuring messages continued to be issued from government during the crisis, it seems probable that the real end to the crisis was simply a consequence of the end of an initial crisis-building phase in which fuel had shifted from underground storage at filling stations to the tanks of concerned motorists. At the beginning of the crisis one might assume that the average driver has a part-full tank, e.g. somewhere around 40 per cent full. By the beginning of the end of the crisis, however, the average driver has a nearly full tank, e.g. 85 per cent full. At that point demand from consumers falls away as absorptive capacity is diminished. If during the crisis the supply chain sources fuel at a steady rate, then the closing phase of the crisis will emerge as supplies to petrol station forecourts starts to exceed demand and during that period stocks in filling station storage tanks will be replenished. The very closing phase of the crisis can be expected to be a need for the supply chain to reduce output to filling stations owing to an eroding absorptive capacity downstream (at this point of both filling stations and vehicles).We suggest that a complex problem of this type is better suited to agent based approaches than to more analytical methods aiming for closed form mathematical solutions. The model presented here is capable of simulating the central aspects of a fuel crisis and is capable of illustrating the processes in informative ways. It should be noted that the major Easter public holiday in England and Wales fell over the long weekend 6-9 April 2012 and by that point the crisis was largely over. Whether the holiday itself contributed to the easing of difficulties is not easily verified, but it would appear to be only a minor issue against what is otherwise an extremely pure example of a fuel crisis panic.

3.4 The model

In this Section, we describe our model of fuel distribution and consumer behaviour. This comprises a *physical model*, which describes a simplified world, and a set of rules to control the motion of our agents and behaviour of the supply chain; a *social-network overlay*, which controls how agents share their observations of their environment; and a *psychological model*, which determines how agents react to these observations. In constructing the model, we embody each of Drogoul's roles (thematician, modeller and computer scientist) ourselves, using the industry observations from Section 3.1 as the basis for our domain model. As with most agent-based simulations, the goal here is not to replicate the exact behaviour of a real system, but rather to produce the most abstract design model which exhibits certain empirically observed phenomena.

3.4.1 Physical model

The one-dimensional world of our model comprises $N \times M$ cells arranged in a ring; for convenience, each cell is assigned an index x, starting from zero. There are N petrol stations, with tank capacity S, spaced regularly around the ring, in cells $i \times M + \lfloor M/2 \rfloor$ for i in the range $[0, N-1]$; M is chosen to be odd, so that every cell has a unique closest station. Every agent has a home cell, which serves as an anchor point for its movement during the simulation and determines the agent's neighbours in the social network overlay. Each cell is home to D agents; the total agent population is therefore $D \times N \times M$. Figure 3 is a graphical representation of the initial state of the model for the $N = 5$, $M = 3$, $D = 3$ case; we see three agents (represented by blue dots) in each of the 15 cells, with a filling station (represented by a green dot) every three cells.

Each agent has a fuel tank of fixed capacity c chosen uniformly within the range $[C_0, C_1]$, and two randomly chosen destinations d_0 and d_1 within R cells of its home; Figure 4 shows the possible destination locations for an agent with home cell 0, for the $N = 5$, $M = 3$, $R = 5$ case. At the start of the simulation, the agent is moved to d_0, and proceeds towards d_1 at a rate of one cell per unit time, consuming one unit of fuel for each cell traversed. Under normal conditions, when it arrives at one destination, it begins to travel towards the other; if the agent has insufficient remaining fuel to reach the new destination and then to drive to the nearest station, it first drives to the nearest station. On arrival at the station, if there is sufficient fuel

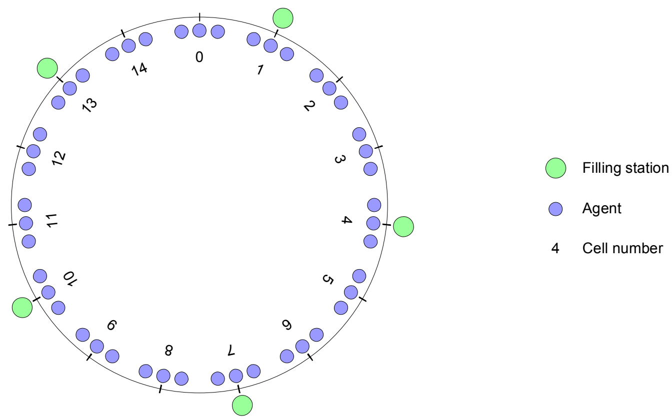

Figure 3: Initial state of model for ($N = 5$, $M = 3$, $D = 3$)

the agent will refill its tank and proceed to the new destination; otherwise it will search adjacent stations in clockwise or counter-clockwise order until it finds one with sufficient fuel, or it has less than M units of fuel remaining, in which case it will join a queue. The preferred direction of search is randomly chosen at the start of each search for each agent and then is maintained until fuel is found. The preferred direction Figure 5 illustrates the unpanicked motion of a typical agent back and forth between d_0 and d_1 (the middle row), with occasional trips to d_0's nearest filling station (the bottom row) or d_1's nearest filling station (the top row) when its fuel supply is nearly exhausted; for simplicity we have omitted the possibility that there is insufficient fuel at the filling station.

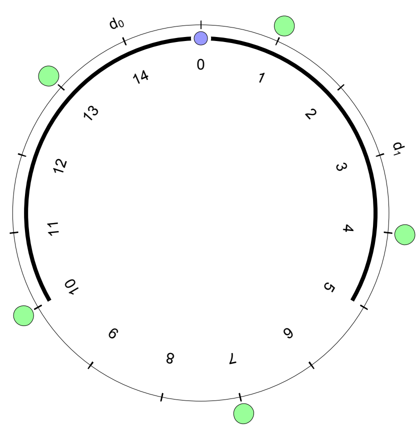

Figure 4: Possible destination locations for an agent with home cell 0 ($N = 5$, $M = 3$, $R = 5$)

Agents do not refuel opportunistically when passing stations, reflecting the finding by Kurani et al [39] that the majority of drivers strongly dislike refuelling their vehicles. To simulate rationing, the model uses functions $U(t)$ and $V(t)$, indicating the maximum and minimum amount of fuel that an agent is permitted to take from a filling station at a given time. In the event that $V(t)$ is non-zero, an agent will not proceed to a filling station unless it can guarantee to accept that much fuel.

Stations are resupplied from a central depot, which for simplicity acts as an infinite-capacity source of fuel. Fuel arrives in tankers with capacity T, and there is a fixed latency of L units of time between a request for fuel and the arrival of the corresponding tanker. Stations do not attempt to forecast demand, and will not

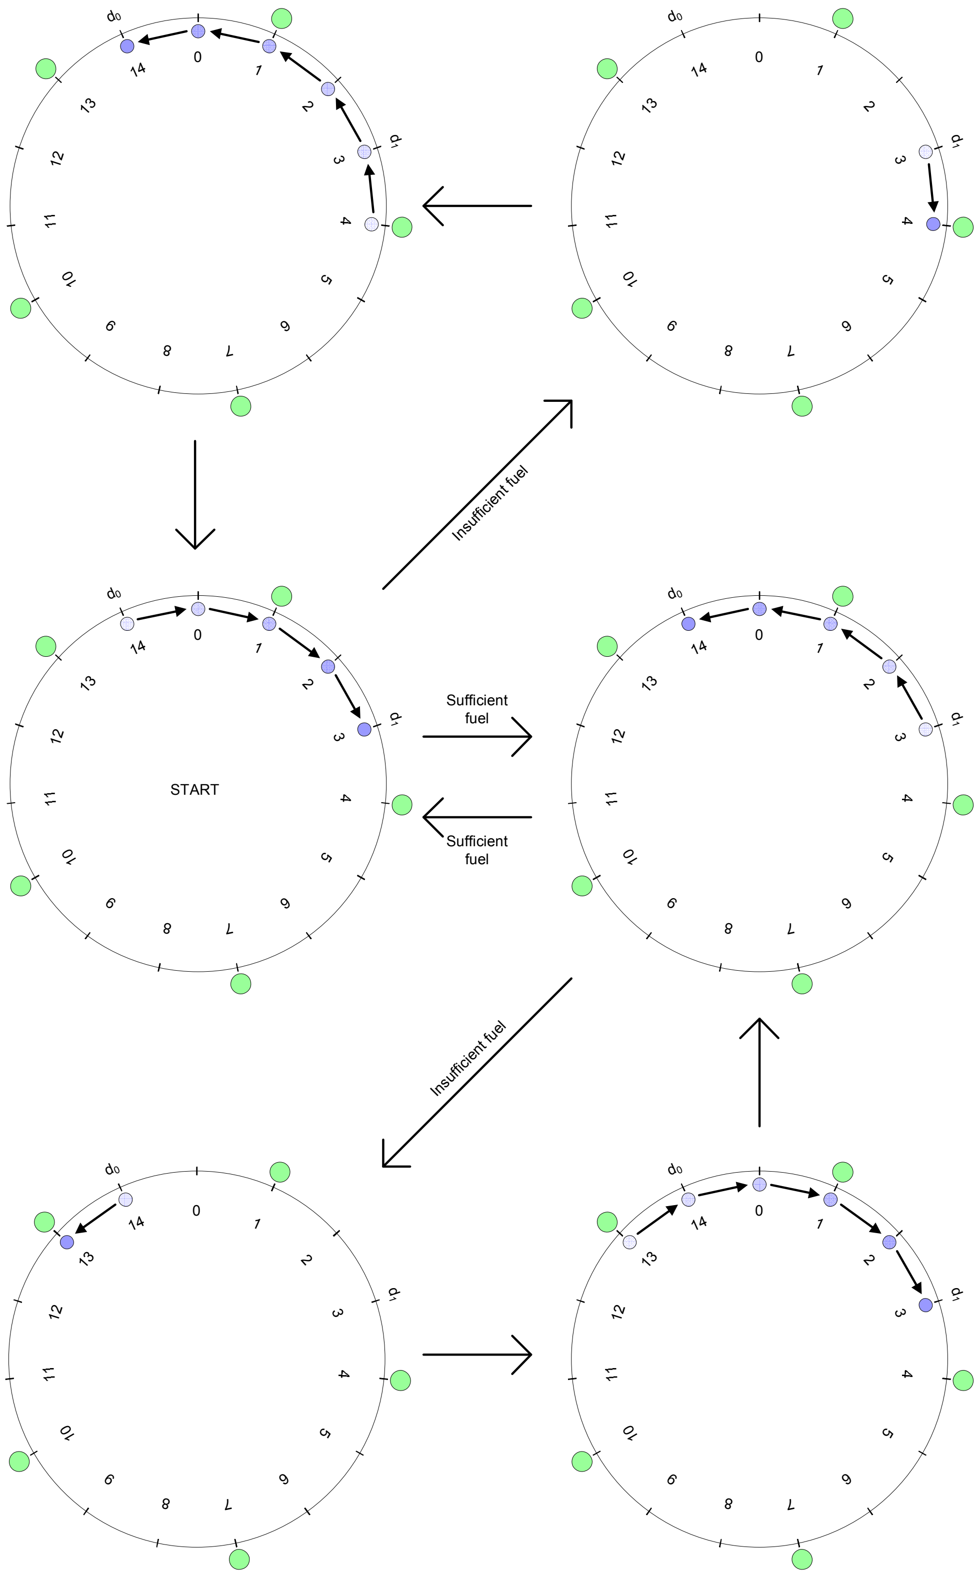

Figure 5: Unpanicked motion of a typical agent

submit a request unless they can guarantee to accept T units of fuel; further requests will be emitted as long as

$$f - e + (i + 1) \times T \leq S$$

where f is the current level of fuel, e is the total fuel requirement of all cars enqueued at the station, and i is the number of outstanding requests which have yet to be fulfilled. Although the trucks are not explicitly modelled as agents in the simulation, the depot implements a scheduling policy to restrict the number of trucks in transit at any given time to a maximum of I, assuming that a truck takes $2L$ units of time to complete a delivery.

Our model has periodic spatial boundary conditions, so no special action is required for agents near the edge of the world. The temporal boundary condition requires more care; introducing the entire population of agents into the simulation at $t = 0$ will result in a large spike in demand as agents refuel near-simultaneously around $t = C$. Each agent's behaviour is deterministic, and in the absence of petrol shortages will repeat with period $p \leq 2C$[1]. We therefore introduce each agent at a time t_0, distributed randomly in the range $[0, p)$; this eliminates spurious demand correlation, bringing the simulation into a steady state by $t = 3C$, after which it is allowed to run until $t = K$.

3.4.2 Social network overlay

To support the flow of information between agents, we construct a scale-free social network overlay using the Watts-Strogatz model. The agents are sorted in order of home cell, and each is connected to the X preceding and X succeeding agents; agents at the beginning of the list are connected to those at the end of the list and vice versa. Each connection is then rewired to a randomly selected agent with probability $0 \leq B \leq 1$, avoiding loops and duplicate connections. Figure 6 shows the formation of the social network for the $N = 3$, $M = 3$, $X = 2$, $B = 0.2$ case; in the graph on the right, 7 of the 36 edges in the original ring lattice have been rewired to random locations.

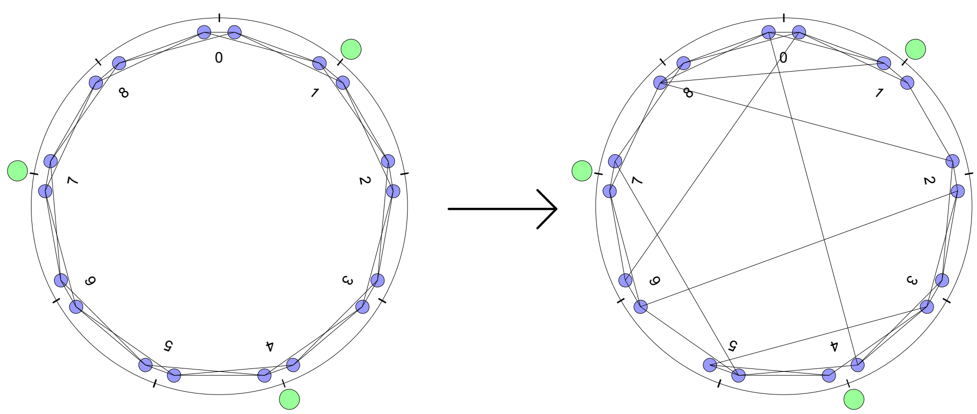

Figure 6: Social network formation ($N = 3$, $M = 3$, $X = 2$, $B = 0.2$)

[1] The upper bound is $2C$ rather than simply C because an agent may alternately refuel from destinations d_0 and d_1.

3.4.3 Psychological model

The key decision facing each agent in our model is whether or not to engage in panic-buying. A panicking agent will attempt a refuelling trip every time it arrives at a destination, rather than waiting until its tank is nearly empty. We utilise a simple stateful psychological model based on Ajzen's theory of planned behaviour to capture this reasoning process.

Each agent has a Boolean state variable $p \in \{True, False\}$, which is $True$ if the agent is currently panicking, or $False$ otherwise. At each time step, the agent has an opportunity to communicate its state to each of its peers in the social network overlay. Communication between any two peers is a random process, occurring with probability P in each time step if the transmitting agent is panicking or with probability Q if it is not; by setting $P \neq Q$ we can model biased communication, in which agents are more likely to communicate one state than the other. An agent can choose to update its state based on direct observation of whether there is a queue at a filling station in the current location, and on communication received from its peers. Figure 7 summarises these sources of data.

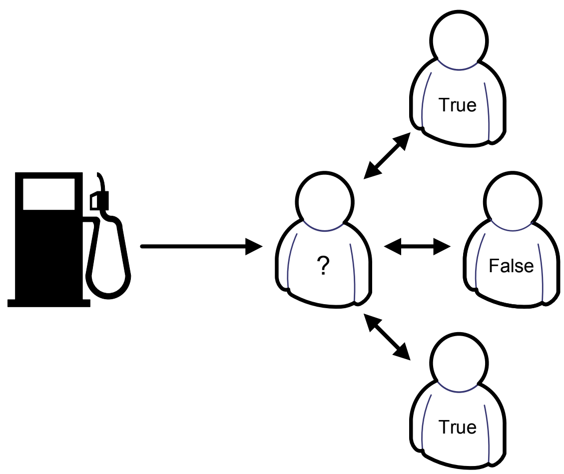

Figure 7: Agents update their state based on communication and direct observation

As noted in Section 2.1, the three antecedents of action in the theory of planned behaviour are attitude, subjective social norms, and perceived behavioural control. For the sake of simplicity we omit the last of these; while more sophisticated simulation might attempt to model an agent's changing beliefs about whether panicking will accomplish anything, we have been able to obtain useful results without incorporating this level of detail. An agent's attitude is controlled directly by the sources of data in 7; each agent maintains a list of the Y most recent communications it has received or observations it has made. If this list contains more negative entries than positive ones, then the agent is predisposed to panic. While an individual's subjective norm relating to panic-buying is by and large socially determined, it may still be expected to vary widely within society, and to remain fixed over the timescale of our simulation. We model an agent's subjective norm by a threshold $1 \leq z \leq Y$; the agent will panic (i.e. set p to $True$) if and only if its queue contains at least z pieces of negative information. In principle at the outset each agent could be assigned a threshold z. In the work presented here the threshold was set to 1 for all agents.

3.4.4 Graphing and reporting

Our simulation provides a range of reporting options, allowing us to examine both aggregate data (fill-level of station tanks, fraction of agents queuing and panicking, number of trucks in use), and local data (proportion of agents with a given home panicking). The former can be depicted using conventional graphs, with the various elements on the y-axis plotted against time on the x-axis. To depict the latter, we use a system of two-dimensional *heat maps*, with distance around the ring on the y-axis and time on the x-axis; each pixel in the map is assigned a colour from the ironbow thermography palette shown in Figure 8, with black indicating no agents at that location panicking and white indicating all agents panicking. Figure 9 shows an example trace and heat map, with a panic event starting at $t = 500$, reaching its peak at around $t = 650$, and ending at $t = 1000$.

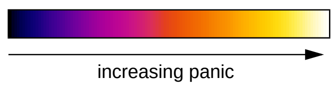

Figure 8: Ironbow thermography palette

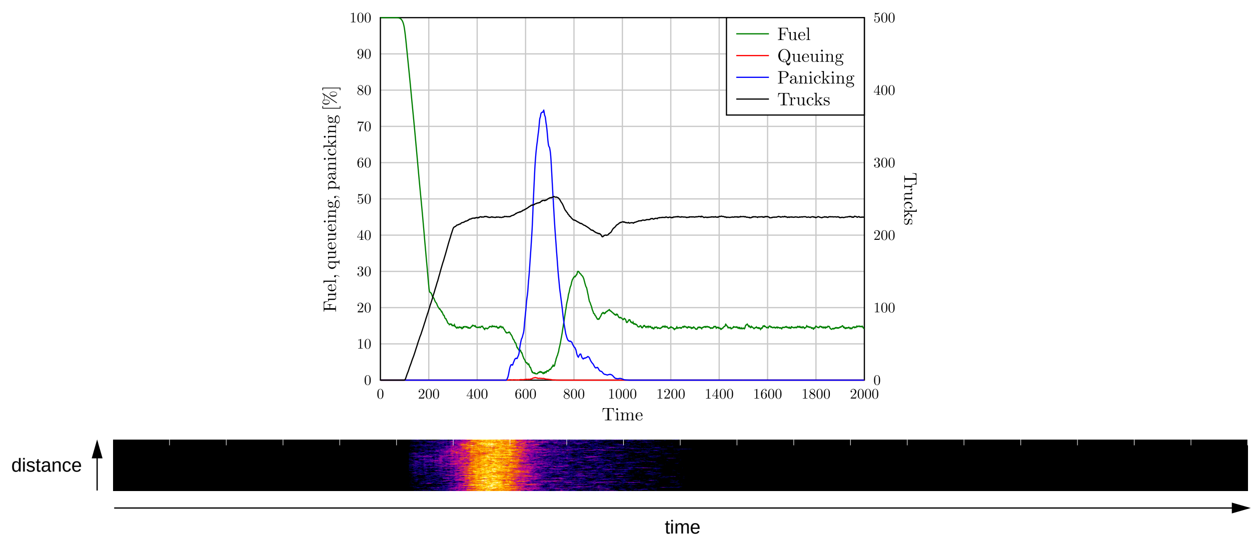

Figure 9: Sample simulation results

For some of our simulations, we provide an aggregate measure of the impact of a disturbance in terms of the number of car-time-units (CTUs) or thousands of car-time-units (kCTUs) spent queuing. One car queuing for 1,000 time units or 1,000 cars queuing for one time unit are judged to have equal impact.

3.4.5 Summary

Table 2 summarises the various parameters of the model.

3.5 Implementation

There exist a large number of software packages which can be used to develop agent-based simulations. Commercial products include AnyLogic [40] and MASS [41]; open-source alternatives such as NetLogo

Parameter	Description
E	random seed
C_0	minimum car tank capacity
C_1	maximum car tank capacity
T	truck tank capacity
S	filling station tank capacity
N	number of filling stations
M	distance between filling stations
D	number of agents at each location
R	range of motion of agents
Y	number of observations to remember
L	latency of fuel delivery
I	maximum number of trucks
K	simulation end time
P	probability of propagating panic
Q	probability of propagating non-panic
X	number of links per agent in social network
B	probability of rewiring link in social network
$U(t)$	maximum ration at time t
$V(t)$	minimum ration at time t

Table 2: Summary of model parameters

[42] are freely available. We have chosen to implement our operational model in the Python programming language [43]. Our choice was motivated by a desire for simulation speed, and a desire easily to permit others to make use oour model and to test our conclusions. We reproduce the source code in full in Appendix A.

We use two techniques to validate our implementation against the design and domain models. By running the simulation for extreme parameter values, we can establish that it has intuitively reasonable behaviour in the limiting case; for example, we check that connecting all agents to each other and disabling the propagation of positive information by setting $P = 1$, $Q = 0$, $B = 0$ and $X = N \times M \times D$ causes all agents to enter the panic state immediately on the first appearance of a queue. We also implement a radically simplified version of the domain model, with very few agents, in a spreadsheet, and check where possible that its results are consistent with those of the primary implementation.

3.6 Experimental results

Having developed and validated our model, we are now in a position to conduct experiments. The goal here is to develop both an empirical understanding of how certain macro-scale features of the real system are determined by the micro-scale behaviour of the participants, and a normative understanding of micro-scale interventions which will produce favourable macro-scale outcomes in the real system. As always with agent-based simulation, our conclusions depend for their validity on the assumption that our design model is a reasonable abstraction of the underlying domain model; we explicitly acknowledge this limitation, framing our conclusions as testable hypotheses about the behaviour of the real world.

We begin by exploring the emergence of oscillations in the supply chain in response to an externally applied demand shock, and develop the concept of *demand coherence*; we show that, under a reasonable set of assumptions about car tank sizes, a common intervention (the addition of “surge” delivery capacity) can exacerbate oscillations by increasing coherence. We then examine the emergence and propagation of panic, highlighting the role of bad-news bias in creating a self-sustaining panic condition, and compare the effectiveness of two different models of censorship. Finally, we evaluate the performance of maximum, minimum and odd-even rationing schemes.

Many of the simulations presented here are variations on a “baseline scenario”, which has 4,500 cars and 10 filling stations on a 90-unit ring. The model parameters for the baseline scenario are shown in Table 3; these are chosen to permit rapid simulation, and to maintain a broadly realistic ratio between aggregate station and car tank capacities. We highlight relevant changes to these parameters in each Section below, and list the full parameter set and procedure for each experiment in Appendix B.

Parameter	Value	Parameter	Value
E	2	L	50
C_0	100	I	1,000
C_1	100	K	2,000
T	2,000	P	0
S	30,000	Q	0
N	10	X	4
M	9	B	0.2
D	50	$U(t)$	100
R	20	$V(t)$	0
Y	1		

Table 3: Model parameters for baseline scenario

3.6.1 Panic-buying can create oscillations through demand coherence

Supply chains are known to be susceptible to oscillatory behaviour when subjected to sudden changes in demand. A recent example is the “Lehman Wave”, first identified by Peels et al [44]; this is a damped, cyclical variation in inventory levels and demand with a period of 12 to 18 months, which has been observed in many supply chains following the bankruptcy of Lehman Brothers in September 2008. In the Lehman case, a rapid, synchronised downward shift in target inventory levels throughout the supply chain caused alternating phases of stocking and destocking; we might expect to see similar dynamics in the fuel supply chain as target fuel levels suddenly increase at the onset of panic, and then decrease to normal levels as the panic ends. There is some evidence, presented in Figure 2, of oscillatory behaviour in the aftermath of the September 2000 fuel protests.

To examine the susceptibility of our model to oscillations, we begin with our baseline scenario and apply an exogenous shock to the system by forcing every agent into the panicked state for 300 units of time starting at $t = 500$. The resulting trace is shown in Figure 10.

We see that the inventory of fuel drops rapidly with the onset of panic at $t = 500$, as agents switch from their normal pattern of infrequent large purchases to a “topping off” behaviour. The number of trucks in the system increases, with the result that inventory returns to its equilibrium level by around $t = 700$. After the panic ends at $t = 800$, however, inventory undergoes two very pronounced oscillations and remains highly

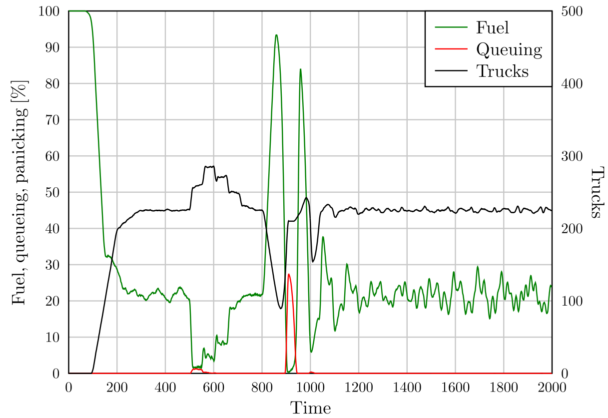

Figure 10: Simulation results ($C_0 = 100$, $C_1 = 100$)

variable for the rest of the simulation. It is notable that during the trough of the first oscillation, nearly 30% of agents are queuing for fuel, many more than at any point during the original disturbance. To understand why these oscillations occur, it is helpful to examine the fraction of agents refueling in each time step, both immediately before and around the trailing edge of the period of panic, as shown in Figure 11.

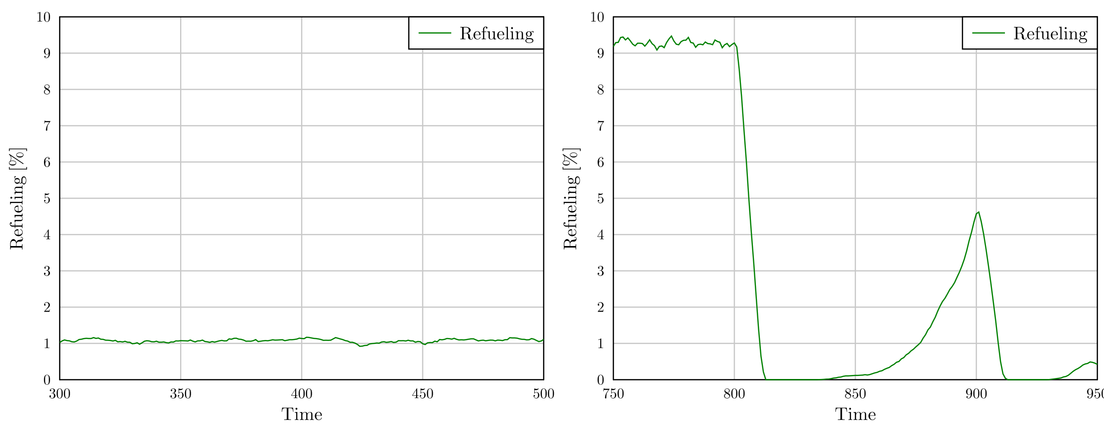

Figure 11: Agents refueling in each time step before and after panic ($C_0 = 100$, $C_1 = 100$)

As we would expect, when the system is in the equilibrium state before the onset of panic, a little over 1% of agents refuel in each time step; cars have 100-unit fuel tanks, and due to the constraints of the physical model must generally refuel shortly before their tank becomes empty. Towards the end of the period of panic, the system has reached another equilibrium state, with nearly 10% of agents refueling in each time step; with agents refueling on average every 11 time units, it follows that the average tank is around 95% full. When the panic ends, there is a brief period during which agents complete refueling trips begun during the panic, followed by a period in which no refueling takes place, as all agents have sufficient fuel. After $t = 850$, the refueling rate increases rapidly, peaking at 4.7% at $t = 900$, as many agents run low on fuel simultaneously; this *coherent demand* drains the fuel inventory completely, and is responsible for the subsequent oscillations in the supply chain.

There are two obvious objections to the applicability of this result to the real world. The first is that occurrences of panic seldom have crisp edges, with the population moving in unison from the non-panicking state to the panicking one or vice versa; we will see in subsequent Sections that oscillations still occur when panic is allowed to spread organically through the social network overlay, rather than being imposed from outside. The second is that real cars have different capacity fuel tanks, and are driven different amounts each day; is it possible that oscillations are purely a by-product of the uniform tank size and consumption rate of our agents? To address this concern, we repeat the previous simulation with fuel tank sizes randomly distributed in the range $[50, 150]$. Figure 12 shows the resulting trace, while Figure 13 shows the fraction of agents refueling in each time step around the trailing edge of the period of panic.

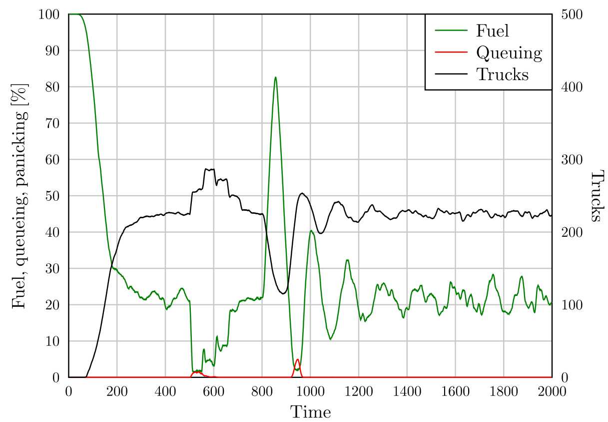

Figure 12: Simulation results ($C_0 = 50$, $C_1 = 150$)

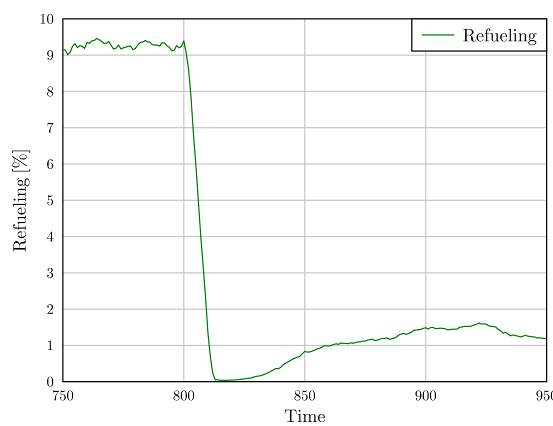

Figure 13: Agents refueling in each time step after panic ($C_0 = 50$, $C_1 = 150$)

We see that oscillations still occur; while they are less pronounced than in the previous case, the peak number of cars queuing during the first trough is still higher than at any point during the original disturbance. Demand is much less coherent, but the refueling rate still rises to 50% above the equilibrium rate at $t = 920$.

3.6.2 Adding delivery capacity can promote oscillation

A common reaction to the appearance of shortages in a supply chain is to bring reserve delivery capacity to bear. At the conclusion of the September 2000 fuel protests, for example, military and spare commercial tankers were used to boost deliveries to 135% of normal levels for a period of several days, with the aim of rapidly restocking depleted filling stations [32].

In its equilibrium state, our baseline model requires between 220 and 225 fuel trucks to be in circulation at any given time. To understand the possible effects of accelerated restocking behaviour, we run the model twice, once with an unconstrained supply of trucks, and once with the number of trucks in circulation constrained to be fewer than 230. In each case, we apply an exogenous panic to every agent beginning at $t = 500$ for 100 units of time. The resulting traces and maps are shown in Figure 14.

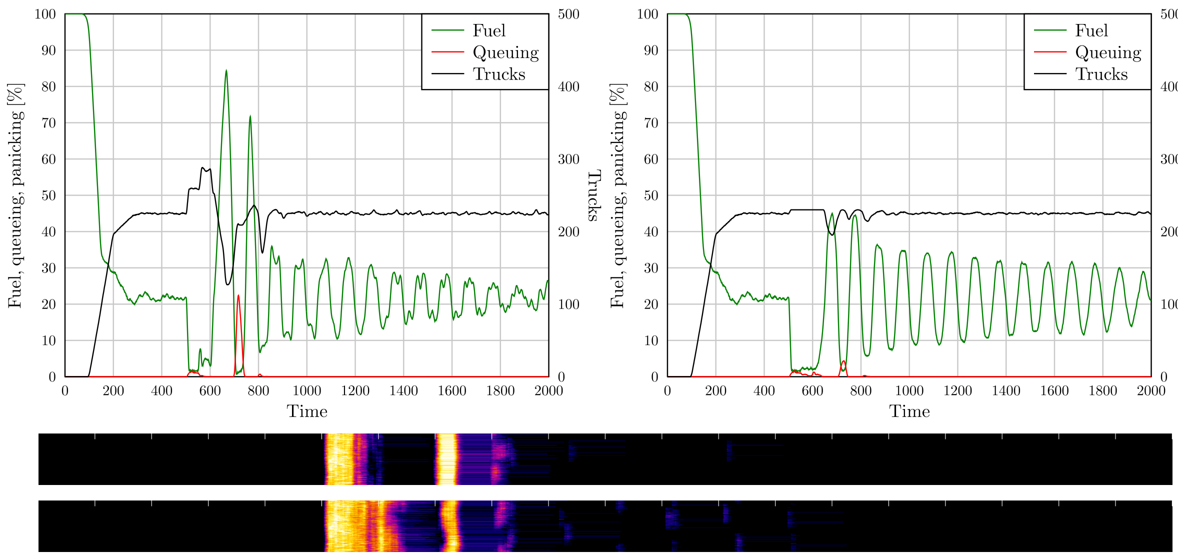

Figure 14: Simulation results without (left and top) and with (right and bottom) cap on truck numbers

In the unconstrained case, the number of trucks climbs rapidly to a peak of 292 as agents throughout the system panic, and filling stations run short of fuel. This increase in supply brings the initial shortage to an end within 100 units of time; because all agents are now able to top off, demand becomes highly coherent, resulting in a larger shortage when many agents refuel at around $t = 700$. In the constrained case, the initial shortage lasts longer, but the secondary shortage is less intense due to reduced coherence; both effects can be clearly seen from the maps. The overall impact of the disturbance is 9.2kCTUs with the constraint, versus 25.9kCTUs without; 22.8kCTUs (88%) of the impact in the unconstrained case occurs during the second shortage. We conclude that adding delivery capacity under shortage conditions can cause damage by increasing demand coherence; the very intervention practised at the end of the September 2000 fuel protests may have contributed directly to the observed oscillations in the fuel supply.

As in the previous Section, it is reasonable to ask whether the observed benefit of restricting the number of trucks is merely a by-product of the uniform tank size and consumption rate of our agents. To address this concern, we repeat the pair of simulations several times, with fuel tank sizes randomly distributed in the range $[100 - a, 100 + a]$ for a in the range $[0, 55]$. Figure 15 shows the increase in impact due to removing the constraint, plotted against $2a$ (the width of the tank-size distribution).

We see that even for a comparatively broad spread of tank sizes there is a substantial benefit to constraining the number of trucks in circulation; at no point does imposing the constraint lead to an increase in impact.

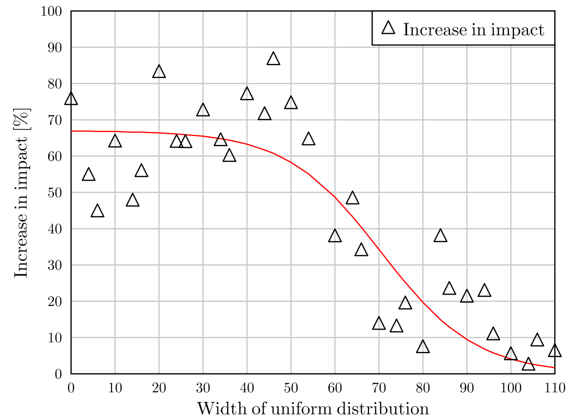

Figure 15: Increase in impact due to adding extra trucks (sigmoid fit)

3.6.3 Unbiased communication can inhibit panic

Supply chains often experience small-scale disruptions, brought about by random variations in demand, or by the failure of a localised element of the distribution system. During the first few days of the September 2000 fuel protests, for example, blockades were sporadic, and regular deliveries continued in most parts of the country; that there was not in any real sense a nationwide crisis encouraged the government to adopt a "business as usual" response to the protests. In these circumstances, communication between individuals might have the effect of spreading panic-buying like a contagion, with consumers inside the affected areas infecting those outside; alternatively it might reassure those inside that the shortages are a local phenomenon, thereby inhibiting panic.

We say that communication is *unbiased* if it has an equal chance of propagating good or bad news between a pair of agents; in the context of our model this requires us to select equal values for the parameters P and Q. To quantify the impact of unbiased communication we run the baseline model twice, once with $P = Q = 0$, denoting "no communication", and once with $P = Q = 0.1$, denoting "unbiased communication"; in each case we simulate a local failure at $t = 500$ by postponing deliveries to one filling station by 15 units of time. The resulting traces and maps are shown in Figure 16.

From the maps, we see that unbiased communication results in low-intensity panic spreading rapidly throughout the system, without ever reaching the level required to trigger a collapse in the supply chain. Without communication, panic initially remains localised in the vicinity of the disturbance, eventually reaching a level where adjacent filling stations are drained of fuel; the resulting disturbance propagates throughout the system, and persists for some hundreds of time units. The overall impact of the disturbance is 514CTUs for unbiased communication, versus 10.2kCTUs for no communication. We conclude that unbiased communication between agents can act to inhibit panic in our model.

This ability to both model and visualise the spatial and temporal evolution of demand is a key advantage of agent-based simulation over alternative approaches to supply-chain modelling such as system dynamics, which typically aggregate inventory levels at a coarser level of granularity [2]. As an example, Figure 17 shows that in the "no communication" case, panic levels at two points on opposite sides of the ring are approximately in antiphase; this level of detail is difficult to obtain using other modelling methodologies.

[2]It is, of course, possible to construct a system dynamics model with sufficient resolution; however, such a model is to all intents and purposes an agent-based simulation.

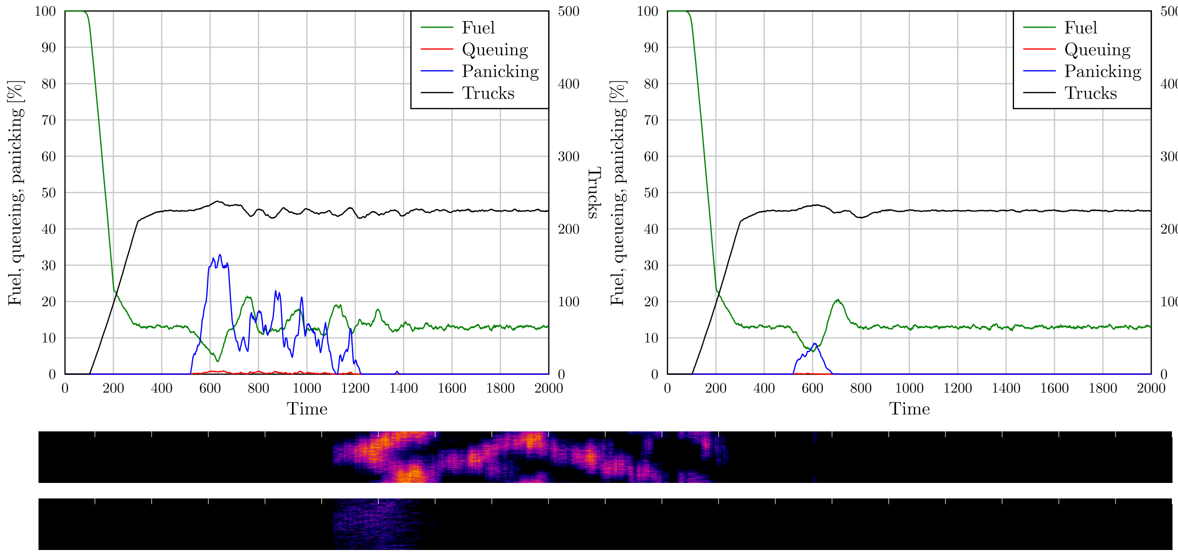

Figure 16: Simulation results without (left and top) and with (right and bottom) unbiased communication

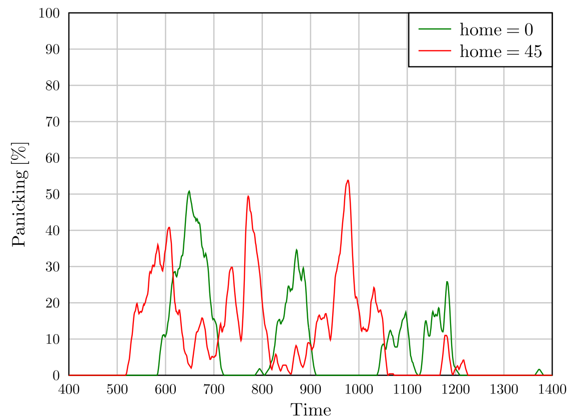

Figure 17: Panic levels at two points on the ring

3.6.4 Even slightly biased communication can promote panic

Although we believe that unbiased communication between individuals can inhibit panic, real-world communication is seldom unbiased. Patterson [45] characterises the tendency of media organisations to focus on negative information as a “bad news bias”, citing several examples from the US political arena. Based on a study of the relative popularity of online New York Times articles, Berger and Milkman [46] report that individuals are more likely to pass on “content that evokes . . . negative (anger or anxiety) emotions characterised by high arousal”. The recent controversy in the UK concerning the MMR vaccine demonstrates how these biases can sustain widespread panic over an extended period in the absence of significant supporting data, and even in the presence of significant data to the contrary.

To illustrate the impact of bad news bias in our model, we begin with the “unbiased communication” configuration described above, and progressively reduce the parameter Q, representing the likelihood of propagating the no-panic state, while holding P constant. As before, we simulate a local failure at $t = 500$ by postponing deliveries. Figure 18 shows maps for progressively larger biases, while Figure 19 shows the variation in impact with bias.

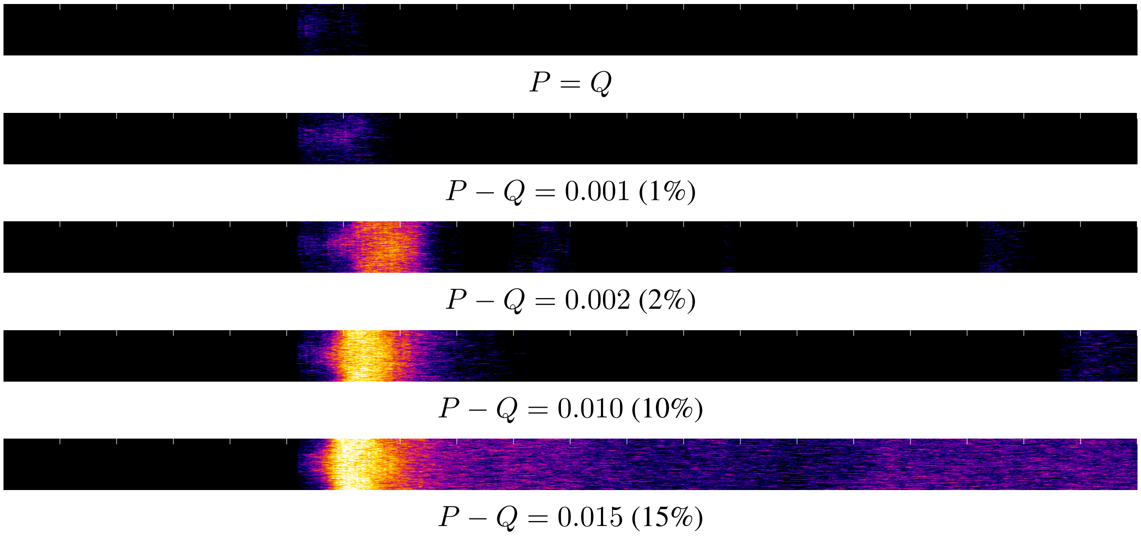

Figure 18: Maps for progressively larger biases

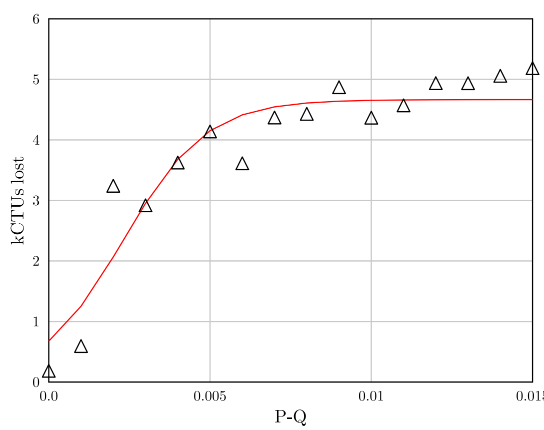

Figure 19: Variation of impact with bias (sigmoid fit)

A bias equal to 1% of the total propagation probability produces a significant increase both in impact and in the incidence of panic, relative to the unbiased case. At a bias of 10%, impact rises to 4.7kCTUs, after which it remains roughly constant. A bias of 15% is sufficient to establish self-sustaining panic within the population of agents; this occurs despite there being no queuing, and therefore no external justification for panic, after $t = 800$. In this last case, an individual agent's positive observations of its environment are overwhelmed by negative information from its peers; this is an example of the "information cascade" phenomenon described by Bikhchandani et al [18].

3.6.5 Censorship is an effective antidote to biased communication

In times of crisis, governments often attempt to control the dissemination of information. In the UK, the DA-Notice (formerly D-Notice) system provides a mechanism by which the government can request that news editors voluntarily refrain from publishing information which is believed to be detrimental to national security; as of February 2011 there are five standing DA-Notices relating to defense and counter-terrorism policy. Although the advent of the internet has limited even authoritarian regimes' ability to control information, it is reasonable to believe that a government might successfully convince or coerce a subset of well-connected individuals to moderate their discussion of an ongoing crisis, with the aim of averting panic.

As noted in Section 2.2, the degree of an agent describes the number of other agents to which it is connected in the social network overlay. Figure 20 shows the degree distribution for our baseline model. The few high-degree agents act as hubs, distributing information between geographically distant regions; we can model the impact of censorship on the incidence of panic by disrupting the flow of information through these agents.

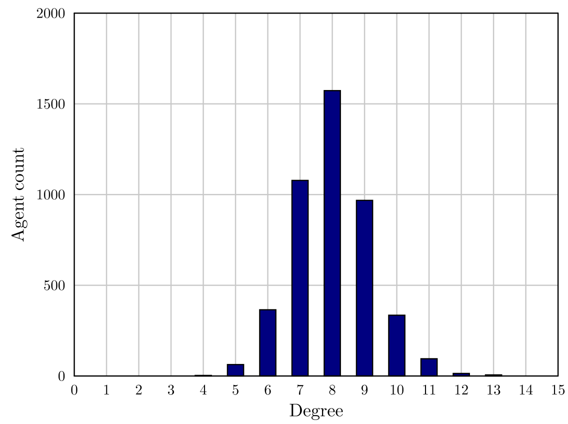

Figure 20: Degree distribution for baseline model

Taking our baseline model with $P = 0.1$ and $Q = 0.085$ (a bias of 15%), we rank the agents by degree, randomly permuting the agents which have the same degree, and censor the highest-ranked agents. If an agent has a lot of connections, it is an influential, high-ranked agent. We pick the n most influential agents, and censor any bad news coming from them. That is, if they try to transmit the news that they saw a queue, we silently destroy the message before it arrives at the recipient, while still allowing good news to propagate. In our first set of experiments, we force each censored agent into the no-panic state at each time step; this corresponds to the "hard" censorship case, in which individuals are induced to actively distribute inaccurate positive information. In our second set, we disable propagation of the panic state from the censored agents; this corresponds to the "soft" censorship case, in which individuals merely avoid distributing negative information. Figure 21 shows the variation in impact with the proportion of agents censored for each case.

Impact follows a sigmoid curve, reaching a minimum of around 100CTUs once 10% of agents have been censored in the hard case, or 15% in the soft case. While little benefit is gained for small interventions, both forms of censorship are effective remedies in the case of biased communication, with hard outperforming soft for moderate-sized interventions.

3.6.6 An enforceable maximum or minimum ration can inhibit panic

A common response to a shortage is to introduce a system of rationing. Where a shortage is expected to persist for an extended period, or is predictable in advance, it is possible to establish a formal system for

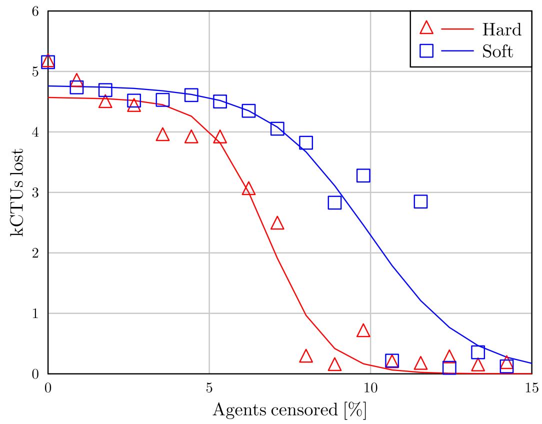

Figure 21: Variation in impact with proportion of agents censored (sigmoid fit)

limiting individual consumption; an example would be the use of ration books in the UK during the Second World War. During the September 2000 fuel protests, an ad hoc system developed [32]: drivers were limited to a certain maximum ration of fuel on each visit to a filling station, with the intent of reducing overall consumption. A maximum ration has the advantage of being easy to implement (by limiting the amount dispensed by each pump), but the disadvantage that the intent is hard to enforce (as a driver can continue to visit filling stations until their tank is full).

To investigate the impact of a maximum ration on our model, we can modify the function $U(t)$, restricting the maximum amount of fuel that an agent can extract from a filling station on each visit. We run the baseline model twice, with a 10% communication bias ($P = 0.1$, $Q = 0.09$) and a simulated local failure at $t = 500$: we use the following values for $U(t)$:

$$\begin{aligned} U(t) &= C_0 && && \text{“no ration” case} \\ U(t) &= C_0, && t < 600 && \text{“ration” case} \\ &= C_0/2, && \text{otherwise} && \end{aligned}$$

The resulting traces and maps are shown in Figure 22.

The introduction of rationing at $t = 600$ immediately halts the decline in fuel inventories, and the rise in the number of panicking agents. By $t = 900$, both metrics have returned to their equilibrium levels. The impact of the shortage is 410CTUs, compared with 2.2kCTUs in the unrationed case.

An alternative approach is the minimum ration, under which cars are prohibited from refueling unless they can accept a certain minimum amount of fuel; the intention here is to reduce transient demand by preventing panicking consumers from topping off their tanks. In comparison with the maximum ration, the intent of this approach is more easily enforceable (once a driver has refueled, they must consume at least the ration before visiting another filling station) but the approach itself is less easily implementable (an attendant must carefully inspect each vehicle prior to fuelling). It has been suggested that it is sufficient to force customers to pay for a minimum ration; we reject this on the grounds that the short-run price elasticity of demand for petrol is low (see Dahl and Sterner [35]) and is likely to decline further under crisis conditions.

To investigate the impact of a minimum ration on our model, we can modify the function $V(t)$, controlling the smallest amount of space that an agent can have in its tank when it refuels. We repeat the previous experiment, setting $U(t) = C_0$, and choosing $V(t)$ as follows:

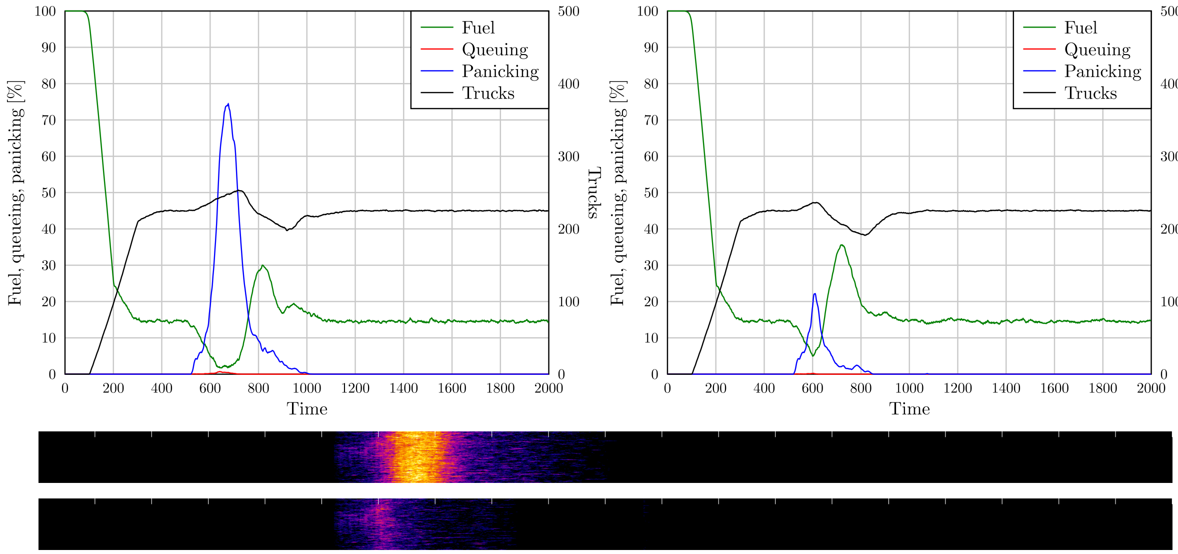

Figure 22: Simulation results without (left and top) and with (right and bottom) maximum ration

$$\begin{aligned} V(t) &= 0 && \text{"no ration" case} \\ V(t) &= 0, \quad t < 600 && \text{"ration" case} \\ &= C_0/2, \quad \text{otherwise} && \end{aligned}$$

The resulting traces and maps are shown in Figure 23.

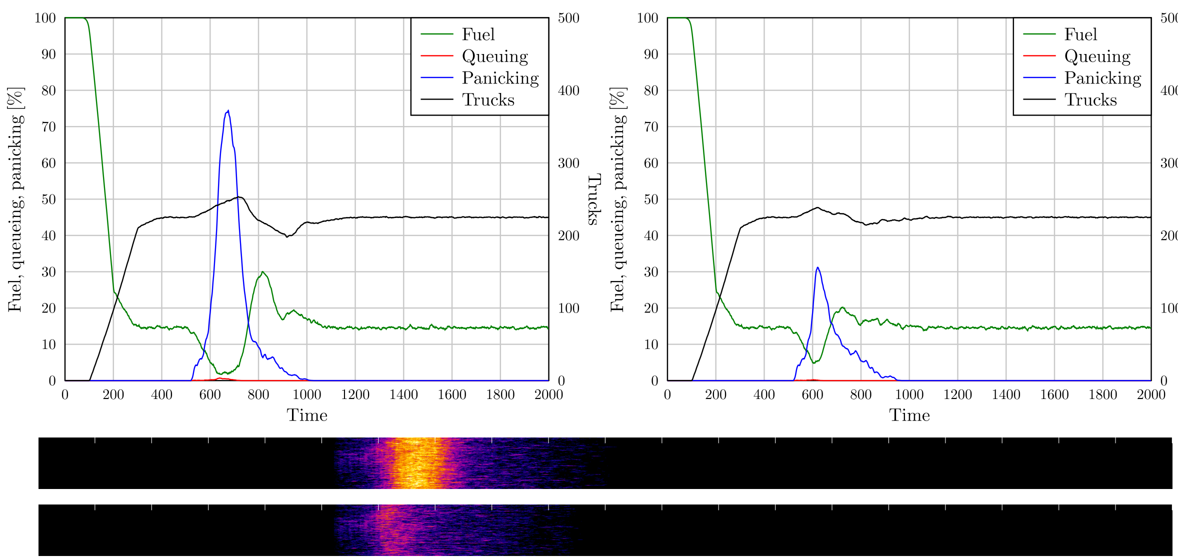

Figure 23: Simulation results without (left and top) and with (right and bottom) minimum ration

Again, the introduction of rationing immediately halts the decline in fuel inventories, and the rise in the number of panicking agents. The system takes slightly longer to return to equilibrium than before, and impact is slightly higher, at 500CTUs.

It is instructive to compare how the two modes of rationing effect a reduction in demand. Figure 24 shows demand decomposed into average purchase size and number of purchases. We see clearly that, after

$t = 500$, panicking agents begin to refuel more often, increasing the purchase rate and decreasing the average purchase size. In the maximum ration case, the reduction in demand after $t = 600$ results from a further rapid drop in average purchase size, as large purchases are completely eliminated. In the minimum ration case, it results from a rapid drop in the number of purchases, as cars with insufficient free tank capacity are eliminated from the pool of potential customers.

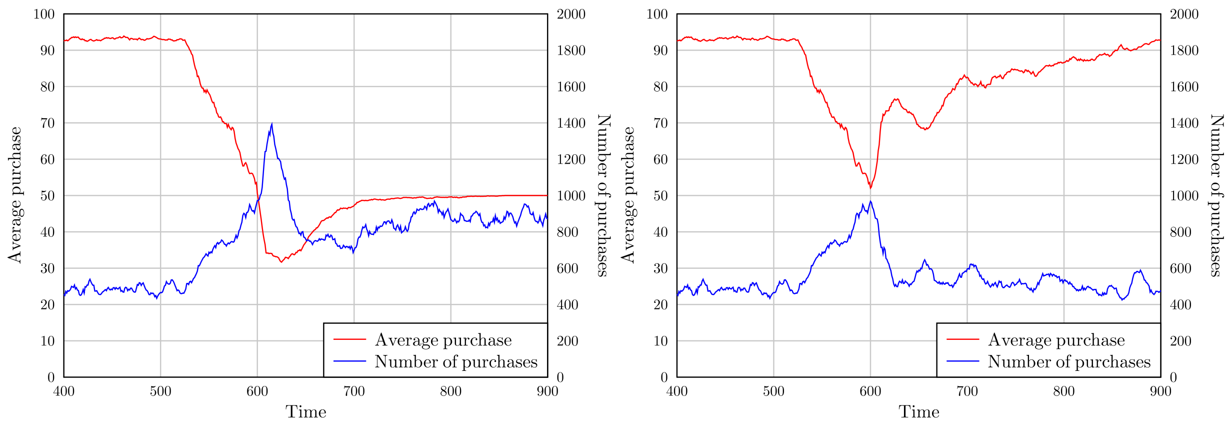

Figure 24: Average purchase and number of purchases for maximum (left) and minimum (right) ration

We conclude that both forms of rationing are similarly effective in inhibiting panic. Choosing between the two requires a judgment about their respective implementability and enforceability.

3.6.7 Odd-even rationing can promote panic by shifting purchases

To control air pollution during the 2008 Beijing Olympics, the Chinese government introduced a scheme whereby cars with odd and even licence plate numbers were only permitted to drive in the city on alternate days [47]. This odd-even approach to rationing is a superficially attractive response to shortages, combining ease of implementation (the eligibility check is trivial) with ease of enforcement (circumventing the check requires a driver to obtain a counterfeit licence plate). Metzger and Goldfarb [48] report that odd-even rationing of petrol was a common policy response in many US states to shortages associated with the oil crises of the 1970s; they present an analytical model of its effects, arguing that under certain assumptions about consumer behaviour it can in fact exacerbate queuing.

Incorporating odd-even rationing into our simulations requires changes to the behaviour of the car and filling station agents. A panicking agent will now only travel to a filling station if it expects to be eligible to refuel on arrival; an agent which is not panicking will refuel early if it calculates that it will no longer to be eligible to refuel when its tank is exhausted. Rather than servicing requests for fuel in strict order, a filling station will now only refuel currently eligible agents, keeping ineligible ones queued. Having made these modifications, we run the baseline model with a 10% communication bias ($P = 0.1$, $Q = 0.09$) and a simulated local failure at $t = 500$, first with no rationing, and then with odd-even rationing with a period of 40 time units starting at $t = 600$; rationing is introduced without warning, denying agents the opportunity to pre-adjust their purchasing behaviour. The resulting traces and maps are shown in Figure 25.

Rationing increases both the intensity and duration of panic; total impact increases to 6.7kCTUs, from 2.2kCTUs without rationing. To explain this phenomenon, it is helpful to plot the proportion of the panicking and non-panicking populations refueling in each time unit, as shown in Figure 26. We see that the

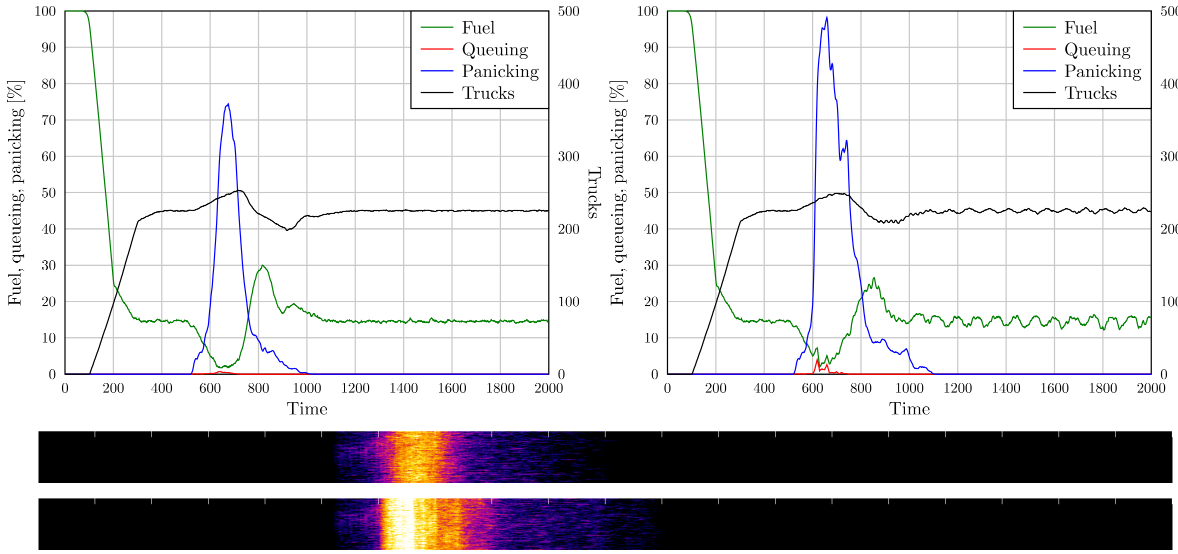

Figure 25: Simulation results without (left and top) and with (right and bottom) odd-even rationing

introduction of rationing at $t = 600$ is accompanied by an abrupt reduction in the proportion of panicking agents refueling; some agents defer their attempts to panic-buy, knowing that they are currently ineligible to purchase fuel. At the same time, we see a marked increase in the proportion of non-panicking agents refueling; these agents have no opportunity to defer, as they only refuel when near exhaustion; some, however, bring their purchase forward, knowing that they will become ineligible to purchase fuel later. In our case, the latter effect outweighs the former, leading to faster consumption, more queuing and more panic.

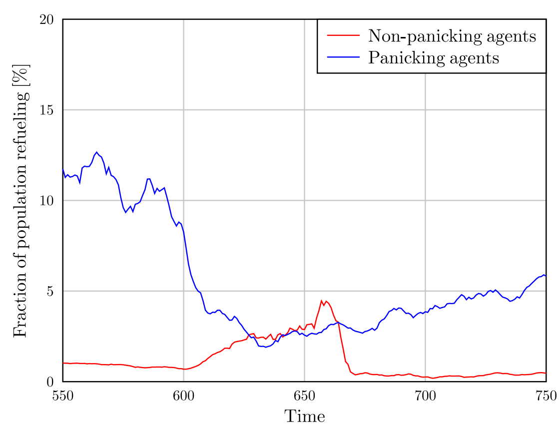

Figure 26: Proportion of panicking and non-panicking population refueling in each time unit

Our results support Metzger and Goldfarb's conclusion that claims for the efficacy of odd-even rationing of petrol rest on "the assumption that infrequent purchasers do not alter their purchasing behaviour in response to [rationing]". Like them, we believe that any benefit is likely to be small, or even negative.

3.7 Policy recommendations

We have drawn a number of conclusions from our simulations, which in turn translate into four recommendations for policymakers confronted with panic-buying:

- **Take steps to limit demand coherence.** When successful, panic-buying leads to a larger than normal degree of correlation between fuel levels in the population of vehicles. Subsequent synchronised refueling can produce aftershocks whose impact is comparable to that of the initial disturbance. While the effect is most pronounced when the panic ends suddenly, or when vehicles have similar natural refueling frequencies, neither is absolutely required. Where possible, governments should attempt to engineer a gradual return to normality; particular care should be taken in using "surge" delivery capacity to restore inventory levels.

- **Promote unbiased communication.** In the case where panic-buying has been triggered by a local failure, unbiased propagation of information about inventory levels can help to reassure consumers in the affected area. Although some consumers outside the affected area may then begin to panic, the effect is likely to be diffuse, and need not necessarily bring about collapse of the entire supply chain. Establishing channels of communication which are immune to allegations of a government "good news bias" will pose a significant challenge.

- **Enforce a maximum or minimum ration.** A maximum ration reduces aggregate demand by cutting the average purchase size; a minimum ration accomplishes the same effect by limiting the pool of eligible purchasers. If introduced during a panic, either can be effective in bringing it to a swift conclusion; the choice between the two should be driven by considerations of implementability and enforceability. Governments may wish to put in place infrastructure to improve the enforceability of a maximum ration or implementability of a minimum ration in advance.

- **Avoid odd-even rationing.** Although it is superficially attractive, this approach encourages some non-panicking individuals to bring their refueling forward in anticipation of being ineligible to refuel when their current fuel supplies are exhausted. Under some circumstances this effect can overwhelm the reduction in the number of currently eligible purchasers, temporarily increasing aggregate demand and promoting panic.

There is evidence that natural bias in communication between individuals can give rise to baseless panic, and that the disruption of communication through high-degree nodes in the social graph is disproportionately effective in limiting its spread. Nonetheless, we stop short of endorsing either "hard" or "soft" censorship as a policy response to the panic-buying of petrol; we do so on ethical grounds, and on the practical basis that, if detected, an attempt at covert censorship is likely to trigger even more widespread panic. Governments may wish to consider censorship during national-security or public-health crises, where ethical considerations carry less weight.

We repeat the caveat that these recommendations depend for their applicability on the assumption that our design model is a reasonable abstraction of the underlying domain model. In applying them, governments should be alert for indications that the macro-scale behaviour of the real system is diverging from the path suggested by our simulations.

3.8 Opportunities to extend this work

The experiments described here involve only small deviations from the baseline scenario described in Section 3.6. There is considerable scope for further experimentation using the model as it stands; in particular it would be possible to examine the effect of:

- **Longer memories and heterogeneous norms.** All of our experiments set $Y = 1$. Increasing this will allow us to model longer agent memories, and to vary social norms with respect to panic across the population.
- **Different social network structures.** Our social network overlay has an average node degree of 8, and a rewiring probability of 20%. It would be possible to quantify the sensitivity of the observed phenomena to these parameters.
- **Different delivery latencies.** All of our experiments set $L = 50$. It is likely that varying the delivery latency will affect the tendency of the model to display oscillatory behaviour.

Each of the elements of our design model is highly simplified in comparison with its domain-model equivalent. Driven by our desire for a system which exhibits empirically observed phenomena, but which is straightforward to understand and experiment upon, our model has in many respects become *simpler* over the course of its development. Key simplifications include:

- A one-dimensional physical environment.
- Uniform spacing and identical capacity of filling stations.
- Homogeneous communication probabilities P and Q.
- Dedicated refueling trips (no opportunism).
- No demand forecasting by filling stations.
- A non-economic behavioural model.

In Sections 3.6, we have presented evidence that our model does indeed exhibit realistic behaviour. Nonetheless, there is scope for further work in verifying the impact of these simplifications by relaxing each one in turn.

Finally, while we remain sceptical of claims for the predictive power of agent-based simulation when applied directly to detailed models of the world, there may be some scope for using simulation to assess the degree of vulnerability of the real UK petroleum supply chain to the effects described in this Paper. Constructing such a model would require a detailed survey of the structure, capacity and latency of the real supply chain, and a much more sophisticated model of consumer psychology and social interaction than has been used here; it is unclear to us whether the benefits of such a model would justify the considerable investment involved.

4 Conclusion

We have demonstrated that agent-based techniques offer a powerful framework for the co-simulation of supply chains and consumers under conditions of transient demand. In the case of panic-buying, even a highly abstract model is capable of reproducing empirically observed phenomena, including oscillations in inventory levels following the restoration of supply, and the rôle of biased communication in sustaining baseless panic. Simulation provides a valuable “sandbox” environment, in which policymakers can experiment with possible responses; we have been able to use our model to evaluate several responses to panic-buying, and to illustrate the mechanisms underlying their success or failure.

In comparison with coarse-grained system dynamics approaches, agent-based simulation is better able to capture spatially-localised phenomena which contribute to the evolution of global demand (*geographical insight*). It also makes explicit the link between the individual and collective behaviour of consumers, and the importance of social interactions as the mechanism by which localised individuals obtain information about the wider state of the world (*social insight*). We suggest that these factors make agent-based simulation the approach of choice in situations where consumers’ access to resources, and personal knowledge of the state of the supply chain, are constrained by geography.

References

[1] M. Gardner, “The fantastic combinations of John Conway’s new solitaire game “life”,” *Scientific American*, vol. 223, pp. 120–123, Oct. 1970.

[2] S. Wolfram, *A New Kind of Science*, 1st ed. Wolfram Media, May 2002. [Online]. Available: http://www.amazon.com/exec/obidos/redirect?tag=citeulike07-20\&path=ASIN/1579550088

[3] R. Axelrod and L. Tesfatsion, “On-line guide for newcomers to agent-based modeling in the social sciences,” Jan. 2011. [Online]. Available: http://www2.econ.iastate.edu/tesfatsi/abmread.htm

[4] T. C. Schelling, “Dynamic models of segregation,” *The Journal of Mathematical Sociology*, vol. 1, no. 2, pp. 143–186, 1971. [Online]. Available: http://dx.doi.org/10.1080/0022250X.1971.9989794

[5] J. M. Epstein, “Modeling civil violence: an agent-based computational approach.” *Proceedings of the National Academy of Sciences of the United States of America*, vol. 99, pp. 7243–7250, 2002. [Online]. Available: http://dx.doi.org/10.1073/pnas.092080199

[6] R. Axelrod, “An Evolutionary Approach to Norms,” *The American Political Science Review*, vol. 80, no. 4, pp. 1095–1111, 1986. [Online]. Available: http://dx.doi.org/10.2307/1960858

[7] J. M. Epstein, “Learning to Be Thoughtless: Social Norms and Individual Computation,” *Computational Economics*, vol. 18, no. 1, pp. 9–24, 2001. [Online]. Available: http://dx.doi.org/10.1023/A:1013810410243

[8] D. J. Hamilton, W. J. Nuttall, and F. A. Roques, “Agent based simulation of technology adoption,” *EPRG Working Paper Series number 0923*, September 2009.

[9] T. Zhang and W. J. Nuttall, “An agent based simulation of smart metering technology adoption,” *EPRG Working Paper Series number 0727*, December 2007.

[10] A. Drogoul, D. Vanbergue, and T. Meurisse, “Multi-agent based simulation: Where are the agents?” in *Proceedings of MABS 2002, Third International Worshop*, 2002, pp. 89–104.

[11] H. Simon, “Bounded rationality in social science: Today and tomorrow,” *Mind & Society*, vol. 1, no. 1, pp. 25–39, 2000. [Online]. Available: http://dx.doi.org/10.1007/BF02512227

[12] E. P. Lazear, “Economic imperialism,” *Quarterly Journal of Economics*, vol. 115, no. 1, pp. 99–146, 2000.

[13] R. H. Thaler, “Toward a positive theory of consumer choice,” *Journal of Economic Behavior & Organization*, vol. 1, no. 1, pp. 39–60, 1980. [Online]. Available: http://dx.doi.org/10.1016/0167-2681(80)90051-7

[14] D. Kahneman, P. Slovic, and A. Tversky, *Judgment under Uncertainty: Heuristics and Biases*, 1st ed. Cambridge University Press, Apr. 1982. [Online]. Available: http://www.amazon.com/exec/obidos/redirect?tag=citeulike07-20\&path=ASIN/0521284147

[15] C. Trepel, C. R. Fox, and R. A. Poldrack, “Prospect theory on the brain? toward a cognitive neuroscience of decision under risk,” *Cognitive Brain Research*, vol. 23, no. 1, pp. 34–50, 2005. [Online]. Available: http://www.sciencedirect.com/science/article/B6SYV-4FMK8F4-3/2/85e1ed063c6f7cdb0b4800c7ebc9239a

[16] R. H. Thaler and C. R. Sunstein, *Nudge: Improving Decisions About Health, Wealth, and Happiness*, updated ed. Penguin (Non-Classics), Feb. 2009.

[17] I. Ajzen, “The theory of planned behavior,” *Organizational Behavior and Human Decision Processes*, vol. 50, no. 2, pp. 179–211, 1991. [Online]. Available: http://www.sciencedirect.com/science/article/B6WP2-4CYG336-DJ/2/9acb128fd8ccf9877db6da912e884f25

[18] S. Bikhchandani, D. Hirshleifer, and I. Welch, “A Theory of Fads, Fashion, Custom, and Cultural Change as Informational Cascades,” *Journal of Political Economy*, vol. 100, no. 5, pp. 992–1026, 1992. [Online]. Available: http://dx.doi.org/10.1086/261849

[19] D. J. de Solla Price, “Networks of Scientific Papers,” *Science*, vol. 149, no. 3683, pp. 510–515, 1965. [Online]. Available: http://dx.doi.org/10.1126/science.149.3683.510

[20] A. L. Barabasi and R. Albert, “Emergence of Scaling in Random Networks,” *Science*, vol. 286, no. 5439, pp. 509–512, 1999.

[21] D. J. Watts and S. H. Strogatz, “Collective dynamics of ‘small-world’ networks,” *Nature*, vol. 393, no. 6684, pp. 440–442, 1998. [Online]. Available: http://dx.doi.org/10.1038/30918

[22] P. Erdős and A. Rényi, “On random graphs I,” *Publicationes Mathematicae*, vol. 6, pp. 290–297, 1959.

[23] UKPIA, “UK Petroleum Industry Association Statistical Review,” June 2010.

[24] F. Boctor, J. Renaud, and F. Cornillier, “Trip packing in petrol stations replenishment,” *The International Journal of Management Science*, vol. 39, no. 1, pp. 86–98, 2011.

[25] S. E. B. Council, “Planning application SE/09/0070 - Erection of petrol filling station and recycling area,” Jan. 2009.

[26] “Profitable detached petrol station plus car wash,” Jan. 2011. [Online]. Available: http://uk.businessesforsale.com/uk/Profitable-Detached-Petrol-Station-Plus-Car-Wash-For-Sale.aspx

[27] “Petrol service station and convenience store,” Jan. 2011. [Online]. Available: http://uk.businessesforsale.com/uk/Petrol-Service-Station-and-Convenience-Store-In-Cardiff-For-Sale.aspx

[28] “Profitable service station in West Kent,” Jan. 2011. [Online]. Available: http://uk.businessesforsale.com/uk/Profitable-Service-Station-in-West-Kent-for-Sale.aspx

[29] UKSA, “Transport statistic bulletin: Vehicle licensing statistics 2009,” 2010.

[30] E. Anderson, “A new model for cycles in retail petrol prices,” *European Journal of Operational Research*, vol. 210, pp. 436–447, 2011.

[31] A. J. Heppenstall, A. J. Evans, M. H. Birkin, and D. O'Sullivan, "The use of hybrid agent based systems to model petrol markets," in *Agent-Based Simulation: From Modeling Methodologies to Real-World Applications*, ser. Agent-Based Social Systems. Springer Tokyo, 2005, vol. 1, pp. 182–193.

[32] N. Robinson, "The politics of the fuel protests: Towards a multi-dimensional explanation," *The Political Quarterly*, vol. 73, no. 1, pp. 58–66, 2002.

[33] A. Babbington, "1bn cost to business of fuel crisis," *The Independent*, 15 September 2000.

[34] P. Hathaway, *The effect of the fuel 'protest' on road traffic*. Department for Environment, Transport and the Regions, 2001.

[35] C. Dahl and T. Sterner, "Analysing gasoline demand elasticities: a survey," *Energy Economics*, vol. 13, no. 3, pp. 203–210, 1991. [Online]. Available: http://www.sciencedirect.com/science/article/B6V7G-458XJF8-3F/2/e505112d4456b0325d75dd8ea75d9a65

[36] P. Walker and D. Milmo, "Army on standby as petrol tanker drivers vote on strike action," *The Guardian*, March 2012. [Online]. Available: http://www.guardian.co.uk/uk/2012/mar/25/army-standby-petrol-tanker-strike

[37] S. Rogers, L. Oliver, and L. Evans, "Uk fuel shortages: the petrol panic mapped," April 2012. [Online]. Available: http://www.guardian.co.uk/news/datablog/interactive/2012/mar/30/uk-fuel-shortages-petrol-panic-map?intcmp=239

[38] SkyNews, "Minister condemned for petrol strike advice," March 2012. [Online]. Available: http://news.sky.com/story/6237/minister-condemned-for-petrol-strike-advice

[39] K. Kurani, D. Sperling, and T. Turrentine, "The marketability of electric vehicles: battery performance and consumer demand for driving range," in *Battery Conference on Applications and Advances, 1996., Eleventh Annual*, 1996, pp. 153–158.

[40] X. Technologies Company, "Software Simulation Tool – AnyLogic," Dec. 2010. [Online]. Available: http://www.xjtek.com

[41] A. International Inc, "The Home of the Multi-Agent Simulation Suite," Jun. 2008. [Online]. Available: http://mass.aitia.ai

[42] U. Wilensky, "NetLogo Home Page," Jun. 2010. [Online]. Available: http://ccl.northwestern.edu/netlogo

[43] Python Software Foundation, "Python programming language – official website," Nov. 2010. [Online]. Available: http://www.python.org

[44] R. Peels, M. Udenio, J. C. Fransoo, M. Wolfs, and T. Hendrikx, "Responding to the Lehman Wave: Sales forecasting and supply management during the credit crisis," *BETA working paper series number 297*, December 2009.

[45] T. E. Patterson, "Bad news, bad governance," *Annals of the American Academy of Political and Social Science*, vol. 546, pp. 97–108, 1996.

[46] J. A. Berger and K. L. Milkman, "Social Transmission, Emotion, and the Virality of Online Content," *Social Science Research Network Working Paper Series*, Dec. 2009.

[47] N. Mulvenney, "Beijing to launch Olympic 'odd-even' car ban," *Reuters News*, 20 June 2008.

[48] M. Metzger and R. Goldfarb, "The effectiveness of odd-even gasoline rationing," *Atlantic Economic Journal*, vol. 11, pp. 49–58, 1983.

A Source code

```
import itertools
import numpy
import random
import sys

from collections import deque

E = 2        # random seed
C0 = 100     # minimum capacity of car fuel tank
C1 = 100     # maximum capacity of car fuel tank
T = 4000     # capacity of truck fuel tank
S = 54000    # capacity of station fuel tank
N = 10       # number of petrol stations
M = 9        # gap between petrol stations
D = 50       # number of cars per location
R = 20       # range of movement of agents
L = 100      # minimum fuel delivery latency
I = 1000     # truck inventory
K = 2000     # number of steps
P = 0        # probability of propagating panic (!!)
Q = 0        # probability of propagating non-panic
X = 4        # number of links per node
B = 0.2      # beta for Watts-Strogatz
A = 10       # averaging window for output
O = 0        # number of agents to censor

G = [None for i in range(K+C1)]

U = [C1 if i < 600 else C1 for i in range(K)]   # maximum ration at each time
V = [0 if i < 600 else 0 for i in range(K)]     # minimum ration at each time

F = [None for i in range(K)]    # force panic vector

SOFT = False
HARD = False

assert O >= -N * M * D and O <= N * M * D

# ironbow palette

PALETTE = [0, 0, 0]  # more RGB data here

def get_color(x):
   x = max(0, min(x, len(PALETTE) / 3 - 1))

   return 0xff000000 | PALETTE[x*3] | PALETTE[x*3+1] << 8 | PALETTE[x*3+2] << 16

# Watts-Strogatz network

def watts_strogatz(num, lim, beta):
   res = list(itertools.chain(*[[(i, (i+j)%num) for j in range(1, lim+1)] for i in range(num)]))

   lookup = set()

   for r in res:
      lookup.add((r[0], r[1]))
      lookup.add((r[1], r[0]))

   for i in range(len(res)):
      if random.random() < beta:
         r = res[i]

         while True:
            j = r[0]
            k = random.randint(0, num-1)

            if j != k and not (j, k) in lookup:
               res[i] = (j, k)

               lookup.remove((r[0], r[1]))
               lookup.remove((r[1], r[0]))
               lookup.add((j, k))
               lookup.add((k, j))
               break

   return res

# utility functions

def can_refuel(u, t):
   if time < 600:
      return True
   else:
      return G[t] == None or G[t] == (u & 1)

def wrap(x):
   return (x + N * M) % (N * M)
```

```
def seek(pos, dest):
   if wrap(dest - pos) < N * M / 2:
      return wrap(pos + 1)
   else:
      return wrap(pos - 1)

def dist(x, y):
   d = wrap(x - y)
   if d < N * M / 2:
      return d
   else:
      return N * M - d

def make_dest(home, exclude):
   while True:
      dest = wrap(home + random.randint(-R, R))

      if dest != exclude and dest != find_station(dest).pos:
         return dest

def find_station(x):
   return stations[x / M]

def calc_period(c, p0, p1):
   STATE_P0 = 0
   STATE_P1 = 1
   STATE_S0 = 2
   STATE_S1 = 3

   s0 = find_station(p0).pos
   s1 = find_station(p1).pos

   p0p1 = dist(p0, p1)
   p0s0 = dist(p0, s0)
   p1s1 = dist(p1, s1)
   s0p1 = dist(s0, p1)
   s1p0 = dist(s1, p0)

   s = STATE_P1
   f = c - s0p1
   b = 0

   while True:
      if s == STATE_P0:
         if f < p0p1 + p1s1:
            f -= p0s0
            s = STATE_S0
         else:
            f -= p0p1
            s = STATE_P1
      elif s == STATE_P1:
         if f < p0p1 + p0s0:
            f -= p1s1
            s = STATE_S1
         else:
            f -= p0p1
            s = STATE_P0
      elif s == STATE_S0:
         return False, b + (c - f)
      elif s == STATE_S1:
         if b == 0:
            b = c - f
            s = STATE_P0
            f = c - s1p0
         else:
            return True, c - f

# depot agent model

def deliver(t, s):
   if t == time:
      s.incoming -= T
      s.fuel += T
      return False
   else:
      return True

class depot:
   def __init__(self):
      self.usage = [0 for i in range(K * 2)]
      self.queue = []

   def request(self, s):
      t = time
      while self.usage[t] == I:
         t += 1

      self.queue.append((t + L, s))

      for i in range(t, t + 2 * L):
         self.usage[i] += 1
```

```
   def postpone(self, s, d):
      self.queue = [(p[0] + d, p[1]) if p[1] == s else p for p in self.queue]

   def tick(self):
      self.queue = [p for p in self.queue if deliver(p[0], p[1])]

# station agent model

REASON_EXHAUSTED = 0
REASON_PANICKING = 1
REASON_NONE      = 2

class station:
   def __init__(self, pos):
      self.pos = pos

      self.incoming = 0
      self.enqueued = 0
      self.capacity = S
      self.fuel = self.capacity

      self.queue = []

      self.taken = [0, 0]
      self.count = [0, 0]

   def enqueue(self, c, reason):
      r = min(U[time], c.capacity - c.fuel)

      self.queue.append((c, r))
      self.enqueued += r

      self.taken[reason] += r
      self.count[reason] += 1

   def tick(self):
      mark = 0

      while mark < len(self.queue):
         c, r = self.queue[mark]

         if can_refuel(c.uid, time):
            if self.fuel < r:
               break
            else:
               self.fuel -= r
               c.fuel += r
               c.state = STATE_WORK
               del self.queue[mark]
               self.enqueued -= r
         else:
            mark += 1

      while self.fuel - self.enqueued + self.incoming + T <= self.capacity:
         d.request(self)
         self.incoming += T

# car agent model

STATE_WORK = 0
STATE_FUEL = 1
STATE_WAIT = 2

g_uid = 0

def last_chance(u, t, f, d0s, d01, d1s):
   while True:
      t += d01
      f -= d01

      if f < d1s:
         return True

      if can_refuel(u, t+d1s):
         return False

      d0s, d1s = d1s, d0s

class car:
   def __init__(self, home):
      self.pos0 = make_dest(home, None)
      self.pos1 = make_dest(home, self.pos0)

      self.capacity = C0 if C0 == C1 else random.randint(C0, C1)
      self.fuel = self.capacity

      start, period = calc_period(self.capacity, self.pos0, self.pos1)

      self.pos = find_station(self.pos1 if start else self.pos0).pos
```

```
      self.fpos = 0
      self.bias = 0

      self.state = STATE_WORK
      self.delay = random.randint(0, period-1)
      self.last = 0

      self.data = {self: False}
      self.conn = []

      self.panic = False

      self.wait = -1

      global g_uid
      self.uid = g_uid
      g_uid += 1

   def tick(self):
      assert self.fuel >= 0

      if time >= self.delay:
         if self.state == STATE_WORK:
            self.pos = seek(self.pos, self.pos1)
            self.fuel -= 1

            if self.pos == self.pos1:
               self.pos0, self.pos1 = self.pos1, self.pos0

               d0s = dist(self.pos0, find_station(self.pos0).pos)
               d01 = dist(self.pos0, self.pos1)
               d1s = dist(self.pos1, find_station(self.pos1).pos)

               if self.capacity - (self.fuel - d0s) >= V[time]:
                  reason = REASON_NONE
                  if last_chance(self.uid, time, self.fuel, d0s, d01, d1s):
                     reason = REASON_EXHAUSTED
                  if F[time] == True or F[time] == None and self.panic and can_refuel(self.uid, time+d0s):
                     reason = REASON_PANICKING

                  if reason != REASON_NONE:
                     self.fpos = find_station(self.pos).pos
                     self.bias = random.randint(0, 1) * 2 - 1
                     self.state = STATE_FUEL
                     self.reason = reason
         elif self.state == STATE_FUEL:
            self.pos = seek(self.pos, self.fpos)
            self.fuel -= 1

            if self.pos == self.fpos:
               station = find_station(self.pos)

               if station.fuel >= station.enqueued + min(U[time], self.capacity - self.fuel) and can_refuel(self.uid, time) or \
                  self.fuel < M:
                  station.enqueue(self, self.reason)
                  self.state = STATE_WAIT
                  self.wait = time
               else:
                  self.fpos = wrap(self.fpos + self.bias * M)

random.seed(E)

# create image

img = numpy.empty((K, N * M), numpy.uint32)
img.shape = N * M, K

# create depot, stations and cars

d = depot()

stations = [station(i * M + M / 2) for i in range(N)]

cars = [car(i / D) for i in range(N * M * D)]

for (i, j) in watts_strogatz(N * M * D, X, B):
   cars[i].conn.append(cars[j])
   cars[j].conn.append(cars[i])

   cars[j].data[cars[i]] = False
   cars[i].data[cars[j]] = False

# build order histogram

histo = {}

for c in cars:
   l = len(c.conn)
   histo.setdefault(l, list()).append(c)

# build censorship list
```

```
for k in histo:
   random.shuffle(histo[k])

censor = list(itertools.chain(*[histo[k] for k in histo]))

if O < 0:
   censor = censor[O:]
else:
   censor = censor[:O]

censor = set(censor)

random.seed(E)

print "fuel,queue,panic,trucks"

sum = 0

for time in range(K):
   if time == 500:
      d.postpone(stations[4], 15)

   for c in cars:
      c.tick()

   for s in stations:
      s.tick()

   d.tick()

   for c in cars:
      c.next = c.panic

      station = find_station(c.pos)

      if c.pos == station.pos:
         c.next = len(station.queue) > 0

      for p in c.conn:
         r = random.random()

         if (p.panic and r < P or not p.panic and r < Q) and not (SOFT and p in censor and p.panic):
            c.next = p.panic

      if HARD and c in censor:
         c.next = False

   for c in cars:
      c.panic = c.next

   count0 = 0
   count1 = 0
   for s in stations:
      count0 += s.fuel
      count1 += len(s.queue)

   count2 = [c.panic for c in cars].count(True)

   count3 = 0
   count4 = 0
   count5 = 0
   count6 = 0
   for s in stations:
      count3 += s.taken[0] + s.taken[1]
      count4 += s.count[0] + s.count[1]
      count5 += s.count[0]
      count6 += s.count[1]
      s.taken = [0, 0]
      s.count = [0, 0]

   count7 = [c.wait == time for c in cars].count(True)

   print count0, ",", count1, ",", count2, ",", d.usage[time], ",", count3, ",", count4, ",", count5, ",", count6, ",", count7

   sum += count1

   for y in range(N * M):
      mycount = 0
      for o in range(D):
         if cars[y*D+o].panic:
            mycount += 1

      img[y,time] = get_color(int(mycount*2.5))

for y in range(10):
   for x in range(99, K, 100):
      img[y,x] = 0xffffffff

from PIL import Image
Image.frombuffer('RGBA', (K, N * M), img, 'raw', 'RGBA', 0, 1).save('my.png')
```

B Model parameters

Here we summarise the model parameters for each of the experiments in Section 3.6. Changes from the baseline scenario are highlighted in red.

Panic-buying can create oscillations through demand coherence

To obtain the "uniform tank size" case, we run the model with the following parameters, applying an exogenous panic to all agents for 300 units beginning at $t = 500$. We run the model again with $C_0 = 50$ and $C_1 = 150$ to obtain the "variable tank size" case.

Parameter	Value	Parameter	Value
E	2	L	50
C_0	100	I	1,000
C_1	100	K	2,000
T	2,000	P	0
S	30,000	Q	0
N	10	X	4
M	9	B	0.2
D	50	$U(t)$	100
R	20	$V(t)$	0
Y	1		

Adding delivery capacity can promote oscillation

To obtain the "constrained" case, we run the model with the following parameters, applying an exogenous panic to all agents for 100 units beginning at $t = 500$. We run the model again with $I = 1000$ to obtain the "unconstrained" case.

Parameter	Value	Parameter	Value
E	2	L	50
C_0	100	I	230
C_1	100	K	2,000
T	2,000	P	0
S	30,000	Q	0
N	10	X	4
M	9	B	0.2
D	50	$U(t)$	100
R	20	$V(t)$	0
Y	1		

Unbiased communication can inhibit panic

To obtain the "communication" case, we run the model with the following parameters, postponing deliveries to the station at $x = 40$ for 15 units at $t = 500$. We run the model again with $P = 0$ and $Q = 0$ to obtain the "no communication" case.

Parameter	Value	Parameter	Value
E	2	L	100
C_0	100	I	1,000
C_1	100	K	2,000
T	4,000	P	0.1
S	54,000	Q	0.1
N	10	X	4
M	9	B	0.2
D	50	$U(t)$	100
R	20	$V(t)$	0
Y	1		

Even slightly biased communication can promote panic

To obtain the "unbiased communication" case, we run the model with the following parameters, postponing deliveries to the station at $x = 40$ for 15 units at $t = 500$. We run the model several times, decrementing Q by 0.001 each time to obtain the "biased communication" cases.

Parameter	Value	Parameter	Value
E	1	L	100
C_0	100	I	1,000
C_1	100	K	2,000
T	4,000	P	0.1
S	54,000	Q	0.1
N	10	X	4
M	9	B	0.2
D	50	$U(t)$	100
R	20	$V(t)$	0
Y	1		

Censorship is an effective antidote to biased communication

To obtain the "no censorship" case, we run the model with the following parameters, postponing deliveries to the station at $x = 40$ for 15 units at $t = 500$. We run the model several times, progressively disabling communication for larger numbers of high-degree agents to obtain the "censorship" cases.

Parameter	Value	Parameter	Value
E	1	L	100
C_0	100	I	1,000
C_1	100	K	2,000
T	4,000	P	0.1
S	54,000	Q	0.085
N	10	X	4
M	9	B	0.2
D	50	$U(t)$	100
R	20	$V(t)$	0
Y	1		

An enforceable maximum or minimum ration can inhibit panic

To obtain the "no ration" case, we run the model with the following parameters, postponing deliveries to the station at $x = 40$ for 15 units at $t = 500$. We then run the model with

$$\begin{aligned} U(t) &= C_0, & t < 600 \\ &= C_0/2, & \text{otherwise} \end{aligned}$$

$$\begin{aligned} V(t) &= 0, & t < 600 \\ &= C_0/2, & \text{otherwise} \end{aligned}$$

to obtain the "maximum ration" and "minimum ration" cases.

Parameter	Value	Parameter	Value
E	2	L	100
C_0	100	I	1,000
C_1	100	K	2,000
T	4,000	P	0.1
S	55,000	Q	0.09
N	10	X	4
M	9	B	0.2
D	50	$U(t)$	100
R	20	$V(t)$	0
Y	1		

Odd-even rationing can promote panic by shifting purchases

To obtain the "no ration" case, we run the model with the following parameters, postponing deliveries to the station at $x = 40$ for 15 units at $t = 500$. We run the model again, prohibiting half the agents from refueling in each 20-unit period to obtain the "ration" case.

Parameter	Value	Parameter	Value
E	2	L	100
C_0	100	I	1,000
C_1	100	K	2,000
T	4,000	P	0.1
S	55,000	Q	0.09
N	10	X	4
M	9	B	0.2
D	50	$U(t)$	100
R	20	$V(t)$	0
Y	1		