

Reculer Pour Mieux Sauter: Resurrecting a Household Budgeting ABM for Deteriorating Economic Conditions

Edmund Chattoe-Brown¹[0000-0001-8232-6896], Corinna Elsenbroich²[0000-0003-1153-4326] and Nigel Gilbert³[0000-0002-5937-2410],

¹ School of Media, Communication and Sociology, University of Leicester, UK
ecb18@le.ac.uk

² MRC/CSO Social and Public Health Sciences Unit, University of Glasgow, UK

³ Department of Sociology, University of Surrey, UK

Abstract. Concern has been expressed in some quarters about the ability of Agent-Based Modelling to *progress* rather than merely proliferate arbitrary models. This paper offers a case study of an elderly Agent-Based Model (hereafter ABM) being resurrected because it seemed suitable to a new use, namely to study the effect of the current UK economic crisis on health and wellbeing. The paper aims to contribute on two levels. One is to discuss the resurrection of the ABM, replicate its original results (enhanced with explanations) and extend it to the new situation. This is very much research in progress. The other aim is to show how, to be progressive, ABMs need to remain in use and effectively accessible to the research community. The case study shows how published ABMs alone are unlikely to satisfy that requirement as time passes (and thus facilitate progressive research). Because this is work in progress, it also attempts to show how some modelling commonplaces (for example that arbitrary models still enhance understanding) actually work in practice. Such aspects are often written out when publishing research that is considered finished.

Keywords: Agent-Based Model, Household Budgeting, Research Methodology, Price Inflation, Wellbeing, NetLogo.

1 Introduction

Concerns have been expressed about the extent to which Agent-Based Modelling constitutes a *progressive* research field rather than simply proliferating “toy” models [1]. Difficulties in replicating and reusing models [2], the uneasy relationship between Agent-Based Modelling and validation to evaluate the empirical performance of ABMs [3] and related aspects of research practice all contribute to justifying this concern. As an effective contrast to this situation, so we can see what is at issue, consider the way that regression models can progress by adding new variables, collecting more (or better) data, incorporating new theoretical insights and so on. Whatever the limitations of this approach to understanding social systems, it provides both an agreed framework for research and a metric for comparing the relative success of models. This paper offers a case study of trying to resurrect an elderly ABM of household money management and

apply it to a new problem, the effect of the current UK economic crisis on health and wellbeing. Most of the paper is devoted to the normal business of describing the ABM and presenting results but there is also some reflection on how not keeping the model in use and available to the wider academic community has made this resurrection more difficult and what it might take for Agent-Based Modelling to shift to a paradigm of making models more robust and extending them rather than starting again from scratch on each occasion (which is what almost invariably happens at present).

The first section of the paper briefly reports the impact of the original ABM. The second section describes the challenges of its resurrection and some results. The third section extends the ABM in a very basic way to the new area of economic crisis and health. The final section reflects on what can be learned from this kind of exercise and where the authors intend to take the research now.

2 Progressive Research: Funding/Publication Contingencies

The ABM resurrected here was the first project undertaken by the first author after his DPhil. With hindsight, the quality of the idea seems considerably better than the practical realisation at that time. In addition, as if often the case, the writing up was compromised by the need to secure further funding so the research done was better and went further than the published record it left behind. Thus one reason for trying to keep models that are innovative (or address distinctive research areas) in use by the ABM community is to avoid worthwhile ideas being too vulnerable to this sort of contingency. In fact, the ABM presented here was only written up in unpublished conference proceedings [4] and this seems likely to have severely limited its subsequent citation. (It should also be recalled that, at that time, the readership for Agent-Based Modelling was much smaller and more scattered, before the advent of JASSS for example.) It also meant that two related empirical and theoretical articles [5, 6], which could have been useful in other attempts to model household budgeting (since the data was collected with model building in mind and the theoretical approach was process based to suit Agent-Based Modelling) had no reason to be seen by, for example, JASSS readers (or indeed by Agent-Based Modellers in general). Investigation using Google Scholar (since this is not intended to be a systematic literature review) shows that citation of any of the project research outputs falls into two categories. Either the ABM is cited by other ABMs to make general methodological points or report its “results” without evaluation or the research is used to make general (non-modelling) points in the wider household budgeting community (again without evaluating the status of the ABM and its outputs). What is not seen is any attempt to develop this ABM or use it as a jumping off point for a *better* model of the same topic. It is in *this* sense that Agent-Based Modelling can be said not to progress. Looking back at this research now, it is very easy to see how the ABM could and should have been better. (Of course, it is inevitable in a case study that other factors might be also be relevant. For example, maybe the ABM just didn’t strike other researchers as any good but critique and building on a flawed model still constitutes progress. Only by looking at more cases could we really get a sense of

whether modelling potential is being systematically lost but there isn't really room to do this in adequate detail in a single paper that must report the ABM as well.)

We now consider the resurrected ABM and the social phenomenon it analyses.

3 Understanding Household Budgeting Using an ABM

The idea of household budgeting is a commonplace. If your television breaks can you just go and buy a replacement or do you have to wait or decide what to forego (and how best to forego it) in order to do so? This simple example highlights two key issues that inspired the original budgeting ABM. Firstly, even now with the development of so called behavioural economics [7], the great majority of models of financial decision making are *attributed* to budgeters rather than being established from their world views as qualitative or experimental data. This is because economics methodologically disapproves of the qualitative methods that are most suitable to establish how households budget in practice and therefore almost never collects this data [8]. This means that specific economic models can be shown to perform poorly against data but not in a way which shows that the whole project of starting from axioms of rational decision might be flawed (if, for example, households budget adaptively or face risk rather than mere uncertainty.) By contrast, qualitative data can directly access the broad "approaches" of households to budgeting: habitual, rational, adaptive and so on [9]. The second strand of the approach, which was over ambitious in hindsight, was that ABM could simply encode the subjective strategies of budgeters gleaned from interviews and explore their implications. However, as we will show, an iterative process broadly along these lines (going back and forth between ABM and data and potentially collecting new data in the light of this process) can still prove valuable.

The broad picture ("approach") revealed by the interviews was not the economic one of laying out all possibilities at once [10], choosing the best and repeating this pattern of consumption infallibly but of making various kinds of "arrangements" to both reduce the uncertainty of household budgets and deal with unavoidable uncertainty. An example of the first kind of arrangement would be putting some small change in a jar by the phone every time you make a phone call to offset the resulting bill. An example of the second kind would be use any windfall income to put a few tins of long lasting food in the cupboard against future trouble. Thus dealing with household expenses is a fundamentally sequential and risky process: What happens if the cooker fails just before pay day?

The core of the original ABM was thus a sequence of objects called "commitments" which could either be income or outgoings: Examples of these commitments would be "you owe rent for the 1st January" or "you got paid on 1st January". Budgeting is thus the task of "processing" these commitments as they arise (with more or less knowledge of their properties including future recurrence) in order to achieve a sustainable lifestyle. At its simplest, processing just involves meeting the commitment (paying the rent) or putting it off (still owing the rent).

But what sustainable means depends on the nature of the commitment (something that traditional economic models of consumption don't address). If you don't pay rent

then you are likely to be evicted after a while and all the money you do not pay, you still owe (although in the *very* long term the debt may be written off once all other avenues are exhausted). If you don't buy food, then you will literally starve to death after a time (though you can also go hungry to cut expenses as poor people regularly do: <https://www.independent.co.uk/news/uk/home-news/cost-of-living-crisis-uk-mother-hospitalised-b2245038.html>). But if you want to spend some real money on V-Bucks (the virtual currency in the computer game Fortnite), there will be no long term consequences if you don't and even the desire to do so may fade away again if unsatisfied (unlike the commitment to pay rent). Interestingly, it actually proved quite hard to find an example good that was harmless if consumed in large quantities and yet also completely optional. In economic models of consumption, the worst possible consequence of your choices is that you are less happy than you might otherwise be. In this ABM, if you buy too many V-Bucks and not enough food, you might die. Thus one way of looking at what the ABM does is to unpack the phenomenon of household budgeting to the point where the strategies that people actually use make more sense in the context of an ecologically valid environment (and are not simply seen as failures of or approximations to abstract rationality which is obliged, for analytical reasons, to assume a totally implausible environment for decision). For example, if commitments are unavoidably risky (and the consequences of not meeting them could be very serious) then people will strongly wish to build up some protective savings if they can.

When the original project was carried out the idea of using qualitative data in ABM was too novel to be really effective (though any project based on the resurrected model would now be able to build on much better work in this field, see for example [11] along with previous experience by the researchers.) Although the interviews provided a lot of insight into practical budgeting in general terms, they didn't offer anything that could just be slotted into an ABM (and we now understand that such an aim may actually not be realistic for various reasons.) Instead, the ABM took a number of budgeting approaches inspired by the interviews (and by common sense) to explore their implications and extend the insight gained by the combination of modelling and data. For practical reasons, the published chapter didn't go far down this path, but it is now possible to go a little further and then apply the ABM to a new problem which has recently arisen. (This ability to react faster may be another reason to keep ABM "in play".) This also illustrates the idea of learning even from models that are known to be incorrect but also points to "units of publication" as a limitation on how much can be achieved at one sitting. As with regression, you can get a lot further in each article or chapter if you don't have to explain the Agent-Based Modelling approach before you start every time.

Consider first the assumption of *no brainer* budgeting. This assumed that you consider all your current commitments and meet any that you can afford. (We assume there is no credit to complicate this picture in the first instance.) It is clear that this is largely *not* what budgeters do but it is useful to understand exactly what happens if you do do it and, perhaps, *why* people therefore do not. This starting point also resonates with another general finding that came from the budgeter interviews, namely that budgeting is a potentially *marginal* activity. If you have too little money for your essential needs, then nothing you do will help. If you have a reasonable amount of money, then you may have limited need to budget in the first place (for example because there are always

adequate savings to smooth over the unexpected). In fact, it is possible that no brainer budgeting actually *does* work provided you have sufficient money and/or your tastes are not “naturally” extravagant - though the ABM would probably still have to be used to get a sense of what sufficient means. (Given its importance, the so-called “non satiation” axiom in consumer theory – that you always want more of something – receives surprisingly anecdotal and hand waving justification.) Thus budgeting is only a phenomenon that makes sense in some middle ground between too little money and enough, where how you arrange your finances actually has the capability to make a difference. ABM explorations can help us understand the scope of this middle ground and how it relates to different environments and budgeting strategies.

The ABM was set up (for comparison of various basic budgeting strategies) with a test set of commitments that is meant to abstract from some important challenges of budgeting. Each simulated month is 30 days long. There has to be an income to “drive” the system. This is 200 “units” and arrives on the first of each month (ticks 0, 30, 60 and so on.) There is also a rent of 80 units which is due on the first of each month too. A “food shop” of 20 units is required every 7 days. Finally, each evening the budgeter has a desire for V-Bucks to game (which cost 5 units). These numbers were chosen so that the budgeter can survive only as long as they buy *no* V-Bucks. In welfare terms they suffer no ill effects from this choice but their budgeting strategy needs to support it (and it makes them unhappy even if it doesn’t harm them). All periodicities and relative prices are almost completely arbitrary (except that, for example, a month’s rent almost certainly *does* cost more than a weekly food shop.) But the aim initially is just to understand how budgeting strategies might work dynamically.

The idea of success in this ABM is already slightly more complicated than just optimising “utility” because of the “realistic” properties of commitments. If you can’t afford food for four weeks, you die of starvation and the simulation run is over. If you can’t pay rent for three months, then you are evicted and although you don’t literally die, you are considered to have “failed” in budgeting and moved outside the scope of this particular ABM (so the simulation run is also over). Only certain commitments (like V-Bucks and food) directly generate “satisfaction” (which is also recorded for comparison of runs). This contradicts the economic view that there is no real distinction between what you want to do and what you need to do. In fact, the reality of this distinction (if you do not do some things then literally “life changing” consequences may follow) underpins a lot of problematic social situations.

The outcome of no brainer budgeting is almost intuitive but not quite. For now, only results from a single budgeter are discussed and the system is deterministic so there is no virtue in reporting repeated simulation runs. Understanding can be achieved by looking at a trace of the code that enumerates all purchases and non purchases of commitments. (The code is written in NetLogo [12] and is available from the first author on request.) Initially, the no brainer budgeter does quite well. Income arrives and, because of the assumed periodicities of commitments given above, they immediately pay rent, buy food and V-Bucks. But the trouble starts on the second day when the budgeter buys V-Bucks (the only thing they currently want). More V-Bucks are bought on the next few days, reducing the money the budgeter holds. There is still enough money left to

pay for a second lot of food but the challenge created by continuing V-Buck consumption worsens and there comes a point where no more food can be afforded in that month. Fortunately, more income is due in the second month thus avoiding the immediate start of starvation but it is here that the no brainer behaviour may not be completely expected (though it is still a strict consequence of the ABM assumptions of course.) In the simplest model, all commitments that are not met persist. This is true of rent in the real world but not straightforwardly true of food and V-Bucks. So come the tick with the second lot of income, food and rent which are really needed to avoid budgetary failure must “take their chances” with all the V-Bucks that could not previously be afforded. Following the logic of the no brainer rule, the budgeter thus spends pretty much all their income on the day it arrives leaving them much less well placed for the coming month than they were for the month before. Whether the budgeter does the things that will save it from starvation or eviction depends on the order in which commitments are added (always newest first) and removed (when they can be afforded based on what else has happened that tick and in previous ticks.) In the second month it is the combination of food that already could not be afforded on the day of the second income and the fact that the budgeter runs out of money so early (and therefore cannot meet *new* food requirements) that means it *does* starve to death within the second month (on tick 50).

What we have learned here is not quite that no brainer budgeting doesn’t work. Instead, having a rigorously defined process has drawn our attention to implications of the ABM that we may wish to amend as implausible. If you desire V-Bucks and regularly can’t buy them for financial reasons, then you may have a bit of a splurge when your income arrives but it seems equally unlikely that you will either just to buy that day’s V-Bucks or try and buy every single lot you missed. (The food case is more complicated and so discussion of it will be deferred.)

It is a commonplace that Agent-Based Modelling allows undemanding (and useful) experimentation and this can easily be demonstrated here. Reducing the price of V-Bucks to 2 units *ceteris paribus* means that the budgeter survives more than twice as long with the same strategy (120 ticks) though they still starve in the end. Similarly, putting income up to 300 (so that “incontinent” V-Bucks purchases are much less harmful to the rest of the budget) means that the budgeter survives still longer (330 ticks) before starving. The ability to change commitments in this way also allows us to explore how effective particular budgeting strategies can be against a range of circumstances and how robust a particular strategy is to changes in circumstances.

We have called this modelling exercise a resurrection. The point is partly to repurpose an ABM that has (possibly undesirably) fallen out of use, partly to be sure that the new ABM reproduces the reported features of the old one (to avoid the general modelling habit of starting from scratch) and partly to address the limitations of the previous study (specifically just reporting *what* happened with different budgeting strategies but not *why* it happened to deepen subsequent understanding and direct model building/data collection.) One experiment with the previous ABM was to keep commitments prioritised (so rent always comes before food and food before V-Bucks.) This was found to have significantly less effect on budgetary success than might have been expected but since there was no explanation forthcoming the value of the result was unclear. Perhaps

it was just a bug? Reproducing this experiment with the new ABM yields a very similar result. Prioritisation only postpones starvation by 7 ticks. Tracing explains what has happened here. On the first income tick, behaviour is unchanged. For most of the first month, prioritisation can do little good because the budgeter only ever wants V-Bucks alone or (rarely) food and V-Bucks. For the few occasions when the latter happens it is only marginally better to prioritise food. Since V-Bucks consumption empties out the bank account as before, for a significant part of the month there is no financial gain for prioritisation to achieve. When the next lot of income arrives, the budgeter behaves slightly less foolishly in paying rent and food first but they still buy all the previously unpurchased V-Bucks they can afford and that empties out the account for the coming month. This reveals another behavioural anomaly in the existing assumptions which is that prioritising food doesn't actually gain you anything when you can't meet subsequent food commitments. (The ABM as defined assesses you not by how many food commitments you do meet but by how many you *don't*.)

This draws attention to a complication of food commitments. For rent, you owe everything you don't pay whatever happens. For V-Bucks, because not buying them has no harmful effects, it doesn't matter if your whim for them vanishes again later. But food is not quite in either category. If you buy just a week's food, then once you eat it all you start to starve. But you can't eat a fortnight of food in a week to postpone starvation. Eating just doesn't work like that. Food commitments are about having food when you need it physiologically (roughly three times a day) and this means that if you can't afford food what you may be doing is running down cupboard stocks initially before you actually start to starve. In this case, when you are better off you may buy a fortnight of food but not entirely for immediate consumption but rather to refill the cupboard for emergencies as well. Interestingly this matches a finding of the original interviews which is about creating pots of money for specific purposes (like a jar full of coins to pay the phone bill.) Stores of food also count as a pot of assets that can stabilise budgeting and help to ensure there is always enough to eat (a very pressing commitment). But this draws attention to the fact that a better model of budgeting may need more sophisticated views of how different kinds of commitment work in terms of frequency, "deferability", consequences and so on. Economics assumes that food, rent and V-Bucks are all "functionally equivalent" in generating utility while even brief reflection shows that this is not the case (and likely to be an excessive simplification). Thus an important advantage of the ABM is to be able (even in a highly stylised sense) to represent the kind of "process diversity" arising from the differing natures of goods.

Another experiment simultaneously allows us to check the findings of the previous research and better understand the implications of forgetting "whims" for V-Bucks. In this experiment we "flush" duplicate commitments of V-Bucks from the system (so the most you ever want is two lots of V-Bucks in a tick.) This increases survival time significantly (to 120 ticks) but does not postpone starvation indefinitely. Tracing the code for this budgeting strategy variant suggests that the problem here is that despite removing implausible unmet demand for V-Bucks (which distorts behaviour on income days), enough V-Bucks *are* consumed to make each month slightly less survivable than the one before. Even with this kind of flushing of implausible V-Bucks demand, commitments pile up and then are rapidly met on income day but leave the budgeter less well

equipped for the future. The situation proves not to be improved by a *combination* of prioritisation and flushing duplicate V-Bucks demand either. This ability to straightforwardly combine different aspects of budgeting strategy in various ways is another known advantage of Agent-Based Modelling.

The final experiment to check the face validity of the original ABM using its new implementation is to add *anticipation* to the model. In this case (but with implausible perfect information assumed) the budgeter can foresee necessary expenses (food and rent) for different periods ahead and only spends on V-Bucks what is not earmarked for these. Here survival times shoot up when the budgeter has as little as 14-day foresight (i. e. not the full length of time it takes between income commitments) and they make it through a full year of budgeting (and still have sustainable income and outgoings.) This quite dramatic result was also found in the original version of the ABM.

Thus the new code (written in NetLogo rather than the pretty much extinct Mac Common LISP of the original chapter) reproduces the main findings of the ABM reported in the original chapter but also, through traces, offers explanations that were previously lacking. These explanations in turn provide a guide to assumptions in the simpler model that may need to be replaced or refined thus allowing for further ABM development (in tandem with the interview data).

While it may be grandiose to call these single budgeter systems complex (though outcomes are definitely path dependent: What you spend now determines what you can spend later), it is fair to say that whether they fail and exactly how they fail is not immediately obvious and it takes the trace function to provide proper explanation and confirm that the results are not simply bugs. Behaviour on income days, when the budgeter has no money and when they only want V-Bucks (or other things too) appear to interact in non-obvious ways that merit further exploration. As suggested, but now also demonstrated, even an obviously wrong model can therefore have ongoing heuristic value.

Now, however, showing that the ABM has been convincingly resurrected and that analysis of its results provides some insight into money management, we turn to its value in understanding the current UK economic crisis and its health impacts.

4 Extending the Model to an Economic Crisis

In discussion, the authors realised that the old ABM might be used in a new way. Instead of looking at the sustainability of budgetary rules in fixed environments, it could explore (just in a very stylised manner at this stage) how budgeters might respond to environmental changes (and particularly crises.) Further the different approach to commitments and budgetary “success” from traditional consumer theory suggests how the ABM could be linked effectively to a major policy concern, namely the health effects of poverty [13]. Once we move away from the economic notion that budgetary decisions are just about how much utility you can achieve, the ABM can start making a contribution to answering questions of the form “If you choose this rather than that, what effect will it have on your long term physical or mental health?” (This doesn’t just apply to crises of course but to other lifestyle decisions such as smoking. Economics

doesn't stress the "life course externality" of consumption choices to the individual but it is an interesting question how prevalent these actually are empirically.) For example, in terms of pure financial responsibility (and approved lack of state involvement) it may make sense for the poor to live on "gruel" and never buy V-Bucks but if doing that for long enough makes you depressed (or merely vitamin deficient) then it may (in some sense) be a false economy. We postpone discussion of some deeper policy implications of this point until the final section since they go beyond mere modelling assumptions.

It is unnecessary to add much to the ABM to show at least a stylised result in this area (linking budgeting to health outcomes). To create a crisis, the simulation just puts the price of food up by one unit per tick probabilistically (rapid "inflation"). It is this that ultimately makes the existing budgetary strategy, whatever it is, unviable.

In the extended ABM (also addressing concerns that may have struck the reader about the assumption that spending on food and rent are completely fixed), it is possible for budgeters to switch from a "standard" food shop to a "cheap" one (and the same for rent). Obviously this cheap food costs less money but it also provides less satisfaction and, by assumption, increases risks of physical illness. The assumptions here are completely stylised but are designed at this stage only to "dock" the ABM with a process for which there is already evidence (<https://commonslibrary.parliament.uk/research-briefings/cbp-9414/>): If budgeters eat badly in the long term then various health conditions are likely to be more prevalent (and the poor will eat badly more often). If they live somewhere unpleasant (particularly if it impacts their socialising which is not yet modelled as a commitment but will be shortly) then it may very well impact their mental health (and again with an obvious inequality dimension). It is assumed that budgeters identify the need to shift to cheap food and housing (and they try cheap food first and only cheap housing if that fails to solve the problem) when their necessary outgoings (i. e. not V-Bucks) exceed their incomings. (This also means, for example, that if the environment changes very rapidly, budgeters may not adapt quickly enough to avoid disaster. In the UK it is not just the size of price changes for utilities like gas that are surprising but also their speed: <https://www.theguardian.com/money/2023/feb/05/energy-prices-to-soar-again-as-jeremy-hunt-rejects-pleas-to-halt-rise>)

Let us start by assuming that the budgeter lives in a very agreeable world of a large income (300 units) and cheap V-Bucks (2 units) and that they are following the no brainer budgeting strategy. Without inflation they can survive the whole run and are actually happy (because they get to buy lots of V-Bucks without harmful consequences for their overall budget). They are also actually saving money (their current account increases rather than staying stable or decreasing) under these assumptions. As we start to increase the rate of inflation (chance per tick of a unit increase in the price of standard food) the accumulation of income slows and then begins to reverse. (The effect on V-Bucks consumption and happiness is not straightforward because although the budgeter doesn't get any smarter about managing their money, their available income for V-Bucks changes in a path dependent way as discussed above). At some point, when inflation reaches a certain level, the budgeter realises that their essentials budget does not "work" any more and switches to cheap food. This solves the problem of financial collapse (and the budgeter therefore does not need to descend to cheap housing, at least for now) and their lifestyle again takes a turn for the better.

There is a lot to learn from this basic experiment (as might be expected from moving the ABM into a new area, however simplistically). Firstly, for straightforwardness, some of the assumptions are flat out unrealistic at this stage suggesting further improvements in the model. For example, cheap food does not suffer inflation. In fact, what we are likely to see is a gradual slide into poverty as first standard food, then cheap food, then standard housing all progressively become unaffordable (https://en.wikipedia.org/wiki/Cycle_of_poverty). In line with the original interview study, we should also expect to see a drawing down of various kinds of resources (savings, credit card debt, food in cupboards) which is a good response to a temporary blip but may actually make the budgeting problem worse if it persists (<https://www.theguardian.com/business/2023/jun/29/uk-households-withdrawing-savings-at-fastest-ever-rate-official-figures-show>). However, the aim is really only to demonstrate the phenomenon at this stage. Secondly, this experiment reveals the interplay between budgeting strategies and the environment. If the price of V-Bucks inflates then this only has harmful consequences for budgeting strategies that do not allow the budgeter to buy less (or no) V-Bucks. By contrast, food inflation affects everyone because food is essential. Thirdly, the ABM presented here allows us to explore (in a very stylised way) a crucial distinction between health and happiness. In the example here, the budgeter goes to cheap food which solves their financial problems and thus allows them to drink V-Bucks again. However cheap food may have objective effects on health (rather than simply being a “matter of taste”.) This opens up an important issue for public policy about the limitations on the “pursuit of happiness” which economic models don’t allow us to discuss sensibly but this ABM does. (The results are not shown here but the ABM assumes a background rate of a generic mental and physical illness – lasting for a certain length of time and having a certain impact on happiness. Eating cheap food is assumed to increase the rate of physical illness and living in cheap accommodation to increase the rates of both physical and mental illness. Thus the ABM is able to quantify the differential health effects of different environments and budgetary strategies. This process is completely unfit for policy at this stage because the process assumptions are completely arbitrary but it shows how, with more research, individual choices and their environment could be linked to outcomes of policy concern through the ABM.) Finally, this example is thought provoking because it does not apply to stereotypically deprived budgeters. Without inflation, they are doing nicely. But one alarming implication of the current financial crisis (and particularly astronomical rises in energy costs) is that if things get sufficiently bad, *anyone* can be vulnerable. As before, we are talking about understanding the *co-variance* of environment, welfare and budgetary strategies rather than assuming that each can be considered in isolation.

5 Discussion and Conclusions

The conclusions of this paper can be divided into three parts. The first is to reflect on the need to keep models “in play” to make ABM effectively progressive. It seems unlikely that anyone would claim that ABM has no role to play in understanding house-

hold budgeting but with the effective disappearance of the original model in the literature, this is the situation that has nonetheless arisen. It has taken more effort and time to get back to a workable ABM in this area (to our knowledge still the only one) and that time could perhaps have been better spent on higher quality (and more timely) policy results from the extended ABM. This point has also been widely made in the context of data and model readiness for the *next* pandemic [14]. Had the model been made available and updated at the “death” of Mac Common LISP, there is at least a chance that others would have taken it up, played with it, taught from it, improved and extended it. It is a commonplace that it is hard to reconstitute old models but this argument gives a new spin on why we would *want* to (or ideally never let them fall out of use in the first place.) This is also important because however industrious and honest researchers are, models are much better if they are engaged with by others “externally” as part of progressive research [15]. For all the authors know there are much better budgeting strategies that should be tested in the ABM than we were able to think of but to discover this takes engagement and debate. (Nonetheless, personal learning is also important. Even going back to an old model as a more experienced researcher provides major gains – for example the need to *explain* – and coincidentally quality assure – results by tracing and not just reporting them.)

The second important dimension of the paper is the policy implications of a more realistic model of household budgeting facilitated by ABM. In a model where people can really die (<https://www.theguardian.com/society/2020/jan/28/disabled-man-starved-to-death-after-dwp-stopped-his-benefits>) it is much more important that policy should stop that happening. The economic approach can make budgeting seem like an inconsequential activity but it is the model and not the reality that is inconsequential. Even if we leave aside the moral dimension (which we should not do) there may be a strictly financial argument for not setting up “false economies” in welfare. Even if people can survive on “gruel” for decades, the indirect health burden this creates (paid for at least in the UK by the NHS) may actually make this an unwise policy to advocate for. (At the same time such long term connections may be deniable and beyond the motivation of governments which are re-elected every few years.) The virtues of individual parsimony are not the end of the story when taking a whole system approach. (There are also issues about whether people know that living on gruel can make you depressed and are thus in a position to decide intelligently as individuals or, even if they know, whether poverty allows them to make the best choices. The poor may simply have to ignore future risks just to ensure present survival. This would be a classic example of a collective action problem for public health over longer time scales than are usually considered.) Finally, although the idea is an old one, this ABM reminds us that the relationship between what we want, what we need and what we like is a complex one which it may take modelling to explore. Simply consuming what we most desire may, in various circumstances, be extremely socially problematic. (And it is important to recall the context that there is already *plenty* of evidence for health inequality: <https://www.health.org.uk/publications/reports/the-marmot-review-10-years-on>).

Thirdly, of course, we contribute to the normal business of modelling. There is now an ABM of household budgeting “in play” with code available to other researchers and results being published (we hope) in more visible places. We intend to develop this

approach further and devote more effort to the model being generally usable. We have not only resurrected a disused model but extended it into an area with significant policy relevance showing how household budgeting and changing environments (like the current inflation crisis) might impact on health. Learning from the earlier project we are now much better placed to iterate modelling, use of existing data and the collection of new data to move the model closer to reality and having some legitimate policy value.

References

1. Bithell, M., Edmonds, B. The systematic comparison of agent-based policy models - It's time we got our act together! *Review of Artificial Societies and Social Simulation*, 11 May, <<https://rofasss.org/2021/05/11/SystComp/>> (2021).
2. Edmonds, B., Hales, D. Replication, replication and replication: Some hard lessons from model alignment. *Journal of Artificial Societies and Social Simulation*, 6(4), <<https://www.jasss.org/6/4/11.html>> (2003).
3. Chattoe-Brown, E. If you want to be cited, don't validate your agent-based model: A tentative hypothesis badly in need of refutation. *Review of Artificial Societies and Social Simulation*, 1 February, <<https://rofasss.org/2022/02/01/citing-od-models/>> (2022).
4. Chattoe, E., Gilbert, N. A simulation of budgetary decision-making based on interview data. Paper presented at *Simulating Societies 95: Approaches to Simulating Social Phenomena and Social Processes*, Boca Raton, FL, <https://www.academia.edu/6722038/A_Simulation_of_Budgetary_Decision_making_Based_on_Interview_Data> (1995).
5. Chattoe, C., and Gilbert, N. Talking about budgets: Time and uncertainty in household decision making. *Sociology*, 33(1), 85-103 (1999).
6. Chattoe, E., Gilbert, N. Understanding consumption: What interviews with retired households can reveal about budgetary decisions. *Sociological Research Online*, 6(3), <<http://www.socresonline.org.uk/6/3/chattoe.html>> (2001).
7. Hargreaves Heap, S. What is the meaning of behavioural economics? *Cambridge Journal of Economics*, 37(5), 985-1000 (2013).
8. Starr, M. Qualitative and mixed-methods research in economics: Surprising growth, promising future. *Journal of Economic Surveys*, 28(2), 238-264 (2014).
9. Janssen, M., Jager, W. An integrated approach to simulating behavioural processes: A case study of the lock-in of consumption patterns. *Journal of Artificial Societies and Social Simulation*, 2(2), <<https://www.jasss.org/2/2/2.html>> (1999).
10. Deaton, A., Muellbauer, J. *Economics and Consumer Behaviour*. Cambridge University Press, Cambridge (1980).
11. Ghorbani, A., Dijkema, G., Schrauwen, N. Structuring qualitative data for agent-based modelling. *Journal of Artificial Societies and Social Simulation*, 18(1), <<https://www.jasss.org/18/1/2.html>> (2015).
12. Wilensky, U. NetLogo. Center for Connected Learning and Computer-Based Modeling, Northwestern University, Evanston, IL, <<http://ccl.northwestern.edu/netlogo/>> (1999).
13. Limb, M. Failure to protect cost of living will increase poverty and health inequalities, warn analysts. *British Medical Journal*, 376, o794 (2022).
14. Gilbert, N., Chattoe-Brown, E., Watts, C., Robertson, D. Why we need more data before the next pandemic. *Sociologica*, 15(3), 125-143 (2021).
15. Edmonds, B. Good modelling takes a lot of time and many eyes. *Review of Artificial Societies and Social Simulation*, 13 April, <<https://rofasss.org/2020/04/13/a-lot-of-time-and-many-eyes/>> (2020).