Economics & Philosophy (2025), 1–23
doi:10.1017/S0266267125000021

ARTICLE

Some problems of causal inference in agent-based macroeconomics

Tobias Henschen

Universität zu Köln, Philosophisches Seminar, Albertus-Magnus-Platz 1, 50923 Köln, Germany
Email: tobias.henschen@uni-koeln.de

(Received 24 November 2022; revised 06 June 2024; accepted 04 November 2024)

Abstract
Guerini and Moneta (2017) have developed a sophisticated method of providing empirical evidence in support of the relations of causal dependence that macroeconomists engaging in agent-based modelling believe obtain in the target system of their models. The paper presents three problems that get in the way of successful applications of this method: problems that have to do with the potential chaos of the target system, the non-measurability of variables standing for individual or aggregate expectations, and the failure of macroeconomic aggregates to screen off individual expectations from the microeconomic quantities that constitute the aggregates. The paper also discusses the in-principle solvability of the three problems and uses a prominent agent-based model (the Keynes + Schumpeter model of the macroeconomy) as a running example.

Keywords: agent-based macroeconomics; causal inference; expectations; chaos; downward causation

1. Introduction

Causal inference, the confirmation of causal hypotheses by citing evidence in their support, is an important branch of macroeconomics. In macroeconomics, causal hypotheses are claims about relations of causal dependence that are believed to obtain between variables standing for macroeconomic aggregates. Causal inference is important because causal hypotheses can be used to justify macroeconomic policy decisions: if Y causally depends on X, then X can be manipulated to change Y, while all other variables in a set of preselected variables remain (largely) unchanged.

Generations of macroeconomists have conducted causal inference in different ways. Representatives of the Cowles Commission (most notably Haavelmo 1944/1995 and Koopmans 1950) used economic theory to derive evidence in support of causal hypotheses: they employed simultaneous equations to model the functional relations between variables standing for macroeconomic aggregates and used economic theory to disentangle causally dependent (endogenous) and causally independent (exogenous) variables.

The Cowles Commission approach was strongly criticized by Lucas (1976) and Sims (1980). The famous Lucas critique states that agents with rational expectations understand the models that the Cowles Commission representatives employed to express causal hypotheses, that these agents anticipate the policy interventions based on these hypotheses (i.e. manipulations of *X*), that they adapt their behaviour accordingly, and that their adapted behaviour counteracts the desired outcome (i.e. the change in *Y*). Sims, by contrast, argued forcefully that many of the theoretical assumptions used to disentangle endogenous and exogenous variables are "incredible".

In response to both criticisms, Kydland and Prescott (1982) developed a "real business cycle" (RBC) model of the economy, which later evolved into the canonical dynamic-stochastic general-equilibrium (DSGE) model, extensions of which are widely used in central banks and other policymaking institutions today. RBC and DSGE models do justice to both criticisms because they model relations of causal dependence that remain invariant to policy interventions, while at the same time relying on the assumption of rational expectations, and because their empirical counterparts – structural vector autoregression (SVAR) models – get along without the distinction of endogenous and exogenous variables.

But RBC and DSGE models and their empirical counterparts are themselves mired in serious problems. The problem with RBC and DSGE models is that they rely on theoretical assumptions that are highly unrealistic: on the assumptions of homogeneous agents who solve optimization problems and form rational expectations, and of economic fluctuations that are generated by exogenous shocks that kick the economy out of equilibrium only temporarily. The problem with SVAR models is that they aren't purely empirical: that very often, restrictions deriving from economic theory need to be imposed to identify the exogenous shocks.[1]

There is an alternative class of models that get along without the unrealistic assumptions of RBC and DSGE models: macroeconomic models of agent-based computational economics (ACE). Macroeconomic ACE models model economic fluctuations as endogenously generated and primarily out-of-equilibrium, and agents as heterogeneous agents who interact directly and form adaptive expectations. ACE models are computational in the sense that their model equations are coded in a structured or object-oriented programming language and run on a computer. Running these models on a computer has the great advantage that it allows for the analysis of large economic fluctuations (including booms and busts) and for the performance of a variety of policy experiments. But a major drawback of ACE models is that their relationship with the empirical evidence is quite unclear.

Researchers engaging in macroeconomic ACE modelling tend to regard a model as "validated" if it is able to reproduce a number of micro- and macroeconomic empirical regularities (so-called "stylized facts"). But the problem with that validation method is that for any empirical regularity, there is an indefinite number of relations of causal dependence that might have generated that regularity. The

[1]These restrictions tend to be just as "incredible" as the theoretical assumptions used to disentangle endogenous and exogenous variables (cf. Cooley and Dwyer 1998).

method of stylized fact reproduction is thus incapable of providing conclusive evidence in support of the causal hypotheses that are meant to justify macroeconomic policy decisions.

Guerini and Moneta (2017) have developed a sophisticated method that is meant to provide precisely that evidence. Their method involves the application of a causal search algorithm that identifies the relations of causal dependence that generate the artificial data, which can be obtained when the model is run on a computer, and the real-world data, which can be obtained from real-world observations of the variables that figure in the equations of the respective ACE model. The model is regarded as validated if the relations of causal dependence that generate the artificial data coincide with the relations of causal dependence that generate the real-world data.

The method developed by Guerini and Moneta (2017) marks an important step forward in the direction of successful causal inference in macroeconomic ACE modelling: they employ a causal search algorithm to identify the parameters of a SVAR without imposing any theoretical restrictions, and they use a similarity measure to assess the extent to which the relations of causal dependence that generate artificial data coincide with the relations of causal dependence that generate real-world data. In the present paper, however, I wish to present three problems that get in the way of causal inference in macroeconomic ACE modelling.

The first problem is that there is at least one class of variables that are likely to act as confounders, and that cannot be included in the model because they cannot be measured. These variables stand for the expectations that agents form about the behaviour of all kinds of variables that matter to them. While macroeconomists engaging in RBC or DSGE modelling model these expectations as rational, macroeconomists engaging in ACE modelling model them as adaptive. And while these expectations are likely to be neither fully rational nor fully adaptive, the problem is that we cannot find out: that we won't be able to decide to what extent expectations are rational or adaptive because we cannot measure them.

The second problem will arise if the target system of macroeconomic ACE models (the macroeconomy) is chaotic. If it is chaotic, it will be sensitive to initial conditions; if it is sensitive to initial conditions, causal inference will be more difficult: all potential confounding (sets of) variables Z that have been omitted from the model will have to retain exactly the same value z over the course of history that the iterated model is supposed to target. If any of these values deviates from z, the sensitive dependence of the target system on initial conditions will imply that the relation of causal dependence that obtains between X and Y according to the model might fail to have a counterpart in reality.

The third problem is that the target system of macroeconomic ACE models incorporates relations of downward causation, and that these relations cannot be shown to obtain if a condition is violated that in the more recent philosophical literature on the causal exclusion argument is referred to as "difference maker sufficiency" (Stern and Eva 2023: section 6). The condition requires essentially that an upper-level variable X, on which a lower-level variable Y_i is believed to causally depend, screen off Y_i from lower-level variables $X_1, X_2, \ldots$ that constitute X (that Y_i be probabilistically independent of $X_1, X_2, \ldots$ given X). Regrettably, that condition is often violated in macroeconomics, and if it is violated, relations of downward causation cannot be shown to obtain.

https://doi.org/10.1017/S0266267125000021 Published online by Cambridge University Press

I will present the three problems in sections 4–6. In section 3, I will analyse the method that Guerini and Moneta (2017) propose to provide empirical evidence in support of causal hypotheses, i.e. in support of hypotheses about the relations of causal dependence that are believed to obtain in the target system of ACE models. In section 2, I will present the core of the ACE model, to which Guerini and Moneta apply their inference method: the core of the Keynes-and-Schumpeter (K+S) model developed by Dosi *et al.* (2008, 2010, 2013, 2015, 2017, 2020). For purposes of illustration, I will often refer to the K+S model. In the concluding section 7, however, I will make it clear that the three problems are likely to impede causal inference in macroeconomics more generally. I will also briefly discuss the steps that I think will have to be taken in each case to bring about an empirical solution to these problems.

2. The (Core of the) K+S Model of the Macroeconomy

The core of the K+S model is the model that Dosi *et al.* (2008) develop to model capital- and consumption-goods production and consumption; to that core Dosi *et al.* (2010) later add a Keynesian component, Dosi *et al.* (2013) Minskyan credit dynamics, and Dosi *et al.* (2015) a monetary policy rule. My presentation of the core will follow the exposition of the "baseline model" that Dosi *et al.* (2017) present in section 2 of their most recent work on the model. But my presentation will skip consumption and instead include one of the extensions of the core: the monetary policy rule.

Dosi *et al.* (2017: 66) model capital-goods production as production of heterogeneous machine tools ("capital goods") by a firm i that uses labour ($L^D_{i,t}$) to produce the tools ($Q_{i,t}$), and that invests a fraction of its past sales to search for process and product innovation ($\vartheta^{in}_{i,t}$) and imitation ($\vartheta^{im}_{i,t}$). While successful process innovation/imitation increases labour productivity ($B_{i,t}$), successful product innovation/imitation impacts the productivity of the machines. Whether the firm successfully innovates/imitates is determined by a draw from a probability distribution, the parameters of which ($\zeta_1, \zeta_2 \in]0;1]$) define the search capabilities of the firm (i.e. its capabilities to innovate/imitate).

The graph in Figure 1 can be used to illustrate capital-goods production. Here, the arrows stand for relations of causal dependence. Dosi *et al.* use causal vocabulary only informally. But they conduct simulated policy experiments that involve manipulations of parameters (for instance, of ζ_1 and ζ_2) and observations of ensuing changes in other variables (for instance, in $Q_{i,t}$) (cf. end of this section). Therefore, the relations that they believe obtain between the variables in capital-goods production qualify as causal in the sense of the interventionist account of causality. The interventionist account also underlies the inference method that Guerini and Moneta (2017) employ when trying to identify the parameters of the K+S model.[2]

[2]Their method is a subtype of the method of discovering conditional independence relations. Woodward (2003: 38) says of that method that it takes as primitive a notion of causal dependence, "the meaning or content" of which is captured by his interventionist account.

https://doi.org/10.1017/S0266267125000021 Published online by Cambridge University Press

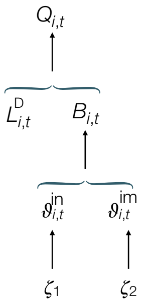

Figure 1. Capital-goods production in the K+S model.

According to the interventionist account, Y causally depends on X relative to a set $\boldsymbol{V}$ of preselected variables iff there is a possible intervention on X that changes Y, while all other variables in $\boldsymbol{V}$ remain unchanged (Woodward 2003: 55, 59). Woodward is not entirely clear about the exact modality of the term 'possible' in 'possible intervention'.[3] But at one point, he characterizes possible interventions as "hypothetical", i.e. as turning the antecedents of counterfactuals into true statements (Hitchcock and Woodward 2003: 183). And in this sense, there is clearly a possible intervention on $L^{D}_{i,t}$ that changes $Q_{i,t}$, while all other variables of the K+S model remain unchanged; there is clearly a possible intervention on $B_{i,t}$ that changes $Q_{i,t}$, while all other variables of the K+S model remain unchanged; and so on. It is quite another question whether there are real interventions: whether the relations of causal dependence that Dosi *et al.* believe obtain in the target system of the K+S model really do obtain in the target system. In order to show that they obtain in the target system, one will have to identify the parameters of the (core of the) K+S model, and sections 4–6 will present three problems that get in the way of parameter identification (or causal inference) in macroeconomics.

Dosi et al. (2010: 1752; 2017: 68) model consumption-goods production as production of final goods: as production of a homogeneous good by firm j whose level of production ($Q_{j,t}$) depends on its desired level of output ($N^{d}_{j,t}$), inventories inherited from the previous period ($N_{j,t\text{-}1}$), and on "adaptive" demand expectations ($D^{e}_{j,t}$), which in turn depend on the demand that the firm faced in the previous two periods ($D_{j,t\text{-}1}$ and $D_{j,t\text{-}2}$), on its demand expectations of the previous period ($D^{e}_{j,t\text{-}1}$), and on the gross domestic product of the previous period ($GDP_{t\text{-}1}$) (cf. Dosi *et al.* 2020: 1494). The graph in Figure 2 can be used to model consumption-goods production. Note that the curved downward arrow stands for a process of downward causation operating from a macroeconomic variable (GDP_t) to a microeconomic variable ($D^{e}_{j,t+1}$).

Dosi *et al.* (2017: 68–69) model price formation as the formation of prices of final (i.e. consumption) goods. Prices of final goods ($p_{j,t}$) equal the unit cost of consumption-goods firm j ($c_{j,t}$) plus a variable mark-up ($\mu_{j,t}$), which depends on the market share of the firm in $t{-}1$ ($f_{j,t\text{-}1}$), which in turn depends on its competitiveness in $t{-}1$ ($E_{j,t\text{-}1}$), which relates directly to $p_{j,t\text{-}1}$ and (inversely) to the possible amount of unfilled demand inherited in $t{-}1$ from the preceding period ($l_{j,t\text{-}1}$). In both sectors,

[3]In one passage, Woodward (2003: 113) rejects the notion of an intervention on variables for which there is no well-defined notion of change (e.g. on variables standing for race, sex or species. In another passage, Woodward (2003: 132) admits interventions that even involve violations of physical laws.

https://doi.org/10.1017/S0266267125000021 Published online by Cambridge University Press

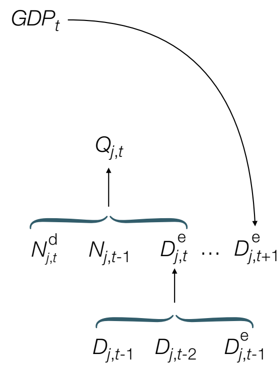

Figure 2. Consumption-goods production in the K+S model.

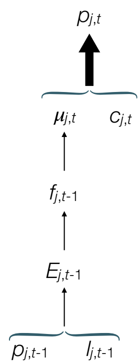

Figure 3. Price formation in the K+S model.

the difference between firm revenues and costs accounts for firms' profits, which change the stock of liquid assets of firms. A firm goes bankrupt if that stock is smaller than 0, or if its market share falls to zero. In both cases, the firm exits the market and is replaced by a new (usually smaller) entrant.

The graph in Figure 3 can be used to illustrate price formation. Here, the fat arrow doesn't denote a relation of causal dependence, but a relation of constitution or asymmetric supervenience. Y asymmetrically supervenes on $X = \{X_1 \ldots X_n\}$ iff changes to Y necessitate changes to X but not vice versa.[4] Price ($p_{j,t}$) clearly supervenes on unit cost ($c_{j,t}$) and mark-up ($\mu_{j,t}$) in this sense: changes to $p_{j,t}$ necessitate changes to $c_{j,t}$ or $\mu_{j,t}$ but changes to $c_{j,t}$ or $\mu_{j,t}$ do not necessitate changes to $p_{j,t}$ (changes to $c_{j,t}$ and $\mu_{j,t}$ can even out exactly). Note that competitiveness ($E_{j,t-1}$), for instance, does not supervene on price ($p_{j,t-1}$) or level of unfilled demand ($l_{j,t-1}$) in this sense: changes to $E_{j,t-1}$ do not necessitate changes to $p_{j,t-1}$ or $l_{j,t-1}$ (they can be due to all kinds of changes, including a change in productivity). Note further that $p_{j,t}$ doesn't causally depend on $c_{j,t}$ or $\mu_{j,t}$. In order for $p_{j,t}$ to causally depend on $c_{j,t}$ or

[4]Cf. Stern and Eva (2021: section 3) for a more thorough analysis of the notion of asymmetric supervenience. I use the terms "to be constituted by" and "to (asymmetrically) supervene on" interchangeably in this paper.

$$C_t + I_t + \Delta N_t = GDP_t$$

$$\sum_{i=1}^{F_1} Q_{i,t} + \sum_{j=1}^{F_2} Q_{j,t}$$

$$Q_{i,t} \qquad Q_{j,t}$$

Figure 4. The constitution of GDP in the K+S model.

$$\pi_t$$

$$CPI_t = \sum_{j=1}^{F_2} p_{j,t}\, f_{j,t} \qquad \pi_t = \frac{CPI_t - CPI_{t-1}}{CPI_{t-1}}$$

$$p_{j,t}$$

Figure 5. The constitution of inflation in the K+S model.

$$\pi_t,\ \pi^T,\ U_t,\ U^T,\ r^T \longrightarrow r_t$$

Figure 6. Monetary policy in the K+S model.

$\mu_{j,t}$, changes to $p_{j,t}$ must not necessitate changes to $c_{j,t}$ or $\mu_{j,t}$.[5] As pointed out above, however, changes to $p_{j,t}$ *do* necessitate changes to $c_{j,t}$ or $\mu_{j,t}$.

The most important relations of supervenience or constitution obtain, of course, between macroeconomic aggregates and the microeconomic quantities that constitute these aggregates. Dosi *et al.* (2017: 70) express GDP, for instance, as the sum of value added of F_1 capital-goods-firms and F_2 consumption-goods-firms and assume that national account identities are satisfied, i.e. that GDP equals aggregate consumption (C_t), investment (I_t) and change in inventories (N_t) (Figure 4). Dosi *et al.* (2015: 171n) define inflation (π_t) as variation of the consumer price index (*CPI*), where the *CPI* is computed as the sum of prices of F_2 consumption-good firms weighted for their market shares (Figure 5).

Finally, according to the monetary policy rule, interest rates (r_t) causally depend on the rate of inflation, the unemployment rate (U_t), and the target rates of interest (r^T), inflation (π^T), and unemployment (U^T) (Figure 6): The policy rule implies that the monetary authority selects the interest rate in response to changes in inflation and unemployment relative to their target levels (cf. Dosi *et al.* 2017: 71).

The graph in Figure 7 illustrates the core of the K+S model plus the monetary policy rule and minus consumption. I should say that the graph doesn't capture every detail of even consumption- and capital-goods production: not every relation of causal dependence or supervenience is covered, and I abstracted from functional form at the micro- and macroeconomic level. But I shall be concerned with problems of causal inference: with problems of providing empirical evidence in support of the relations of causal dependence denoted by the thin arrows. And for the purpose of discussing these problems, the graph is sufficiently detailed.

Before turning to these problems, however, I will have to present the method that Dosi *et al.* (2017: sections 4 and 5) employ to "validate" their model. They validate their model by (1) selecting initial values for variables, by (2) selecting benchmark

[5]The requirement that changes to $p_{j,t}$ do not necessitate changes to $c_{j,t}$ or $\mu_{j,t}$ instantiates a condition that Woodward (2015: 316) refers to as "independent fixability".

https://doi.org/10.1017/S0266267125000021 Published online by Cambridge University Press

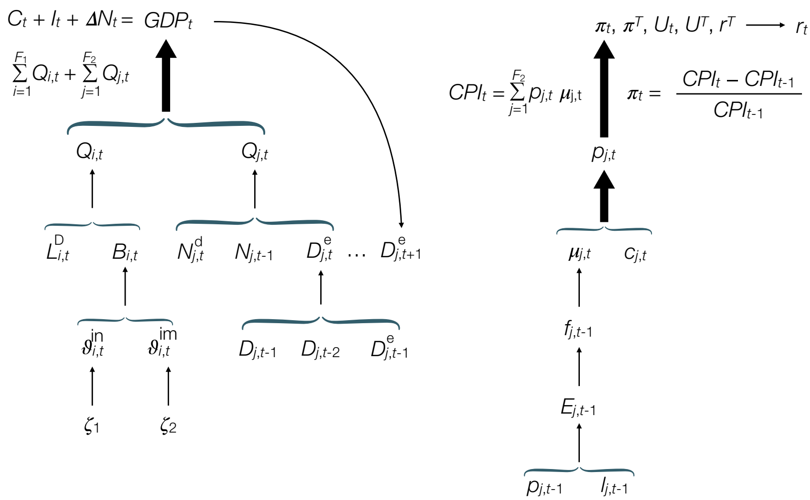

Figure 7. The core of the K+S model plus the monetary policy rule and minus consumption.

parameters, i.e. by calibrating model parameters (by selecting 50 for the number of capital-good firms, 200 for the number of consumption-good firms, 0.30 for $\zeta_{1,2}$, 0.4 for φ, and so on)[6], by (3) coding the model equations in a structured or object-oriented programming language, by (4) running the model on a computer an arbitrary number of times (e.g. 5000 times), and by (5) reproducing a number of micro- and macroeconomic empirical regularities (so-called "stylized facts").

The micro- and macroeconomic stylized facts that Dosi et al. (2008: section 2) say the core of the K+S model is able to reproduce include the heterogeneity of firms in terms of productivity; the fat tails of the GDP growth rate distribution; short-lived recessions; and cross-correlations of macro-variables (more specifically, the pro-cyclicality of consumption, productivity, nominal wages, and inflation, and the counter-cyclicality of unemployment, prices and mark-ups).

Dosi *et al.* (2017: 65) claim that the validated model is "well-equipped to explore . . . the short and long-run effects of economic policies". Among the policies whose effects they explore are innovation policies, Keynesian fiscal policies and policies of income redistribution. When exploring the effects of innovation policies, for instance, they (1) take the validated model as a point of departure, (2) manipulate the parameters measuring the firms' search capabilities (ζ_1 and ζ_2), (3) run the model 100 times, (4) observe that after 100 runs, the GDP growth rate is higher than that of the validated model if $\zeta_{1,2} > 0.3$, and lower if $\zeta_{1,2} < 0.3$, and (5) conclude that "policies favoring innovation promote faster growth" (Dosi *et al.* 2017: 76).

[6]Benchmark parameter selection in macroeconomic ACE modelling is similar to the classical calibration exercise that Kydland and Prescott (1982) introduced in the context of DSGE and SVAR modelling. But a major difference is that in ACE modelling, benchmark parameter selection depends less on mainstream microeconomic (that is, general equilibrium) theory. According to many ACE macroeconomists, that is the reason why ACE models outperform DSGE models in terms of the number of stylized facts they are able to reproduce (cf. e.g. Delli Gatti *et al.* 2008).

https://doi.org/10.1017/S0266267125000021 Published online by Cambridge University Press

3. Causal Inference in Macroeconomic ACE Modelling

The claim that the validated model is "well-equipped" to explore policy effects is precisely the claim that Guerini and Moneta deny. They argue that the method of stylized fact reproduction is not sufficiently rigorous:

> Economics, as any scientific discipline intended to inform policy, has inevitably addressed questions related to identification and measurement of causes and effects. . . . The quality of ABMs has been up to now evaluated according to the ex-post ability in reproducing a number of stylized facts . . . We argue that such an evaluation strategy is not rigorous enough. Indeed the reproduction, no matter how robust, of a set of statistical properties of the data by a model is a quite weak form of validation, since, in general, given a set of statistical dependencies there are possibly many causal structures which may have generated them. Thus models which incorporate different causal structures, on which diverse and even opposite practical policy suggestions can be grounded, may well replicate the same empirical facts. (Guerini and Moneta 2017: section 1)

The method of stylized fact reproduction is not "rigorous enough" because for any stylized fact, there is an indefinite number of relations of causal dependence that might have generated the fact. An example of a stylized fact that the K+S model is able to reproduce is the correlation (and pro-cyclicality) of consumption and inflation (cf. section 2). That correlation is compatible with three relations of causal dependence: consumption causally depends on inflation; inflation causally depends on consumption; or both inflation and consumption depend causally on a third variable (or set of variables). The third relation (with the set of variables acting as a confounder) divides into an indefinite number of relations, depending on the relations of causal dependence that obtain between the variables in the set.

The inference method that Guerini and Moneta (2017) employ to provide empirical evidence in support of the relations of causal dependence between the micro- and macro-variables of the K+S model is the five-step procedure of (1) transforming artificial and real-world data in a way that allows them to be directly comparable, of (2) testing for the assumptions of statistical equilibrium and ergodicity, of (3) estimating the parameters of a reduced-form vector autoregression (VAR) model, of (4) running a specific search algorithm to identify the relations of causal dependence generating the artificial data and those generating the real-world data, and by (5) comparing both sets of relations.

The primary purpose of step (1) is to select k variables of interest and to adjust the time series that provide values for these variables in the artificial and real-world data sets. The assumptions to be tested in step (2) need to be satisfied in order for any model to be able to approximate the data generating process. They require that the observed time series have distributional properties that are time-independent ("statistical equilibrium"), and that they are a random sample of a multivariate stochastic process ("ergodicity"). Note that these requirements don't coincide with the requirement of insensitivity to initial conditions to be discussed in section 5 below.

https://doi.org/10.1017/S0266267125000021 Published online by Cambridge University Press

The VAR model of the artificial and real-world data represents the value of each of the k variables at t as a linear combination of contemporaneous and p lagged values of all k variables:

$$\mathbf{Y_t} = \mathbf{BY_t} + \mathbf{\Gamma_1 Y_{t-1}} + \ldots + \mathbf{\Gamma_p Y_{t-p}} + \boldsymbol{\varepsilon}_\mathbf{t},$$

where $\mathbf{Y_t}$, $\mathbf{Y_{t\text{-}1}}$ … are vectors of the values of variables Y_1 … Y_k at t, $t-1$ … $t-p$, where $\mathbf{B}$ and $\mathbf{\Gamma_1}$ … $\mathbf{\Gamma_p}$ are matrices collecting the model parameters, and where $\boldsymbol{\varepsilon}_\mathbf{t}$ is a vector of error terms. The diagonal elements of $\mathbf{B}$ are equal to zero (no value depends on itself), and the error terms in $\boldsymbol{\varepsilon}_\mathbf{t}$ are assumed to be mutually independent. Subtracting $\mathbf{BY_t}$ from both sides of the equation yields the so-called structural vector autoregression (SVAR) model:

$$\mathbf{\Gamma_0 Y_t} = \mathbf{\Gamma_1 Y_{t-1}} + \ldots + \mathbf{\Gamma_p Y_{t-p}} + \boldsymbol{\varepsilon}_\mathbf{t},$$

where $\mathbf{\Gamma_0} = \mathrm{I} - \mathrm{B}$. Dividing both sides of the SVAR model by $\mathbf{\Gamma_0}$ gives the reduced-form VAR model:

$$\begin{aligned}\mathbf{Y_t} &= \mathbf{\Gamma_0^{-1}\Gamma_1 Y_{t-1}} + \ldots + \mathbf{\Gamma_0^{-1}\Gamma_p Y_{t-p}} + \mathbf{\Gamma_0^{-1}}\boldsymbol{\varepsilon}_\mathbf{t} \\ &= \mathbf{A_1 Y_{t-1}} + \ldots + \mathbf{A_p Y_{t-p}} + \mathbf{u_t}.\end{aligned}$$

The endogeneity of the contemporaneous variables biases estimations of the parameters of the VAR and SVAR models. It is therefore the parameters contained in $\mathbf{A_1}$ … $\mathbf{A_p}$ that are estimated in step (3).

The parameters contained in $\mathbf{A_1}$ … $\mathbf{A_p}$, however, are not the ones that figure in the causal model (i.e. the model representing relations of causal dependence between the k variables). In order to identify the parameters of the causal model, the parameters contained in $\mathbf{\Gamma_0}$, $\mathbf{\Gamma_1}$ … $\mathbf{\Gamma_p}$ need to be recovered. If the parameters in $\mathbf{\Gamma_0}$ are recovered, they can be used (together with the parameters in $\mathbf{A_1}$ … $\mathbf{A_p}$) to identify the parameters in $\mathbf{\Gamma_1}$ … $\mathbf{\Gamma_p}$. But the problem with the parameters in $\mathbf{\Gamma_0}$ is that any inversion of $\mathbf{\Gamma_0}$ is compatible with the parameters contained in $\mathbf{A_1}$ … $\mathbf{A_p}$ (as long as the inverse is unit-diagonal).

Step (4) is meant to solve this problem by running a causal search algorithm: the LiNGAM. The LiNGAM is an algorithm relying on the assumption that the variables in $\boldsymbol{V}$ can be represented as a Linear Non-Gaussian Acyclic Model. Like any other causal search algorithm, the LiNGAM operates by forming a complete undirected graph, and by testing for conditional independence relations to eliminate unnecessary edges and to direct the remaining ones.[7] But other search algorithms result in a (so-called Markov equivalent) class of observationally equivalent directed graphs: they narrow down the class of observationally equivalent equation systems but stop short of fully identifying the parameters of the causal model. The LiNGAM, by contrast, stands out as being capable of identifying the parameters of the causal model completely. Guerini and Moneta run the LiNGAM to discover the causal dependencies among the residuals $\mathbf{u_t} = \mathbf{\Gamma_0^{-1}}\boldsymbol{\varepsilon}_\mathbf{t}$ of the reduced-form VAR model. These causal dependencies give the parameters contained in $\mathbf{\Gamma_0}$, which are then used to identify the parameters in $\mathbf{\Gamma_1}$ … $\mathbf{\Gamma_p}$.

[7]Cf. Glymour *et al.* (2019) for a review of some of the most common search algorithms, their assumptions and output.

https://doi.org/10.1017/S0266267125000021 Published online by Cambridge University Press

In step (5), they use a "similarity measure", which indicates, by what percentage the relations generating the artificial and real-world data coincide with one another.

When applying their five-step procedure to the K+S model, Guerini and Moneta (2017: section 4) select only the 6 macroeconomic variables (GDP_t, C_t, I_t, U_t, π_t and r_t) as variables of interest. They use US time series data from the first quarter of 1959 to the second quarter of 2014 to provide values for these variables, and they adjust the length of the artificial data set accordingly. They test for the assumptions of statistical equilibrium and ergodicity and conclude that both assumptions are "reasonable". They then use the two data sets to specify two VAR models and to estimate the parameters that figure in these models. Next, they employ the LiNGAM to discover the causal dependencies among the residuals $\mathbf{u_t} = \mathbf{\Gamma_0^{-1}}\boldsymbol{\varepsilon}_\mathbf{t}$ of the reduced-form VAR models. Finally, the similarity measure tells them that the percentage by which the relations of causal dependence generating the artificial and real-world data coincide is in the interval between 65% and 80%, depending on the technique that is used to estimate the parameters of the reduced-form VAR models. Guerini and Moneta (2017: 21) say that they "guess that this is a positive result" for the K+S model.

Dosi *et al.* (2020: 1500) agree that this is a positive result: they say that Guerini and Moneta (2017) have "externally validated" the relations of causal dependence "between macroeconomic variables within the K+S model". I am not entirely sure, however, whether this is a positive result. It implies that at least 20% of the relations of causal dependence generating the artificial data have no counterpart in reality, or that at least 20% of the relations of causal dependence generating the real-world data have no counterpart in the model.

I also want to emphasize that the causal inference method that Guerini and Moneta (2017) propose is meant to be an inference method for the K+S model as a whole, and not only for the macro-part of the model. Their method is meant to provide empirical evidence in support of *all* relations of causal dependence that obtain in the target system of the K+S model, and not only in support of the relations of causal dependence that obtain at the macro-level of that system.

I finally want to point out that Guerini and Moneta need to provide empirical evidence in support of the relations of causal dependence that obtain at the micro- *and* macro-level of the system if that evidence is supposed to be evidence in support of the efficacy of the policies that Dosi *et al.* explore. These (innovation, Keynesian fiscal, and income redistribution) policies all involve, after all, manipulations of benchmark parameters at the *micro*-level.

Guerini and Moneta (2017: 5) suggest that there are problems for the application of their method to the K+S model as a whole, and that these problems are practical in nature and not of the in-principle kind: they say that "representing into a unique model every single micro-mechanism at work in a complex economy is a very difficult task". I agree that there are practical problems for the application of their method to the K+S model as a whole. But I also think that in addition to these practical problems, there are in-principle problems for the application of the inference method to the K+S model as a whole: three problems on which the following three sections will elaborate.

4. The Non-measurability of Expectations

The first problem has to do with our inability to measure individual or aggregate expectations. The claim that expectations cannot be measured may come as a surprise to theorists who conduct surveys to measure expectations on a regular basis. But Kevin Hoover (2001: 137) explains why we cannot measure expectations by conducting surveys:

> True, people form expectations and act upon them . . . , but such expectations do not exist independently of the actions they affect; they are not palpable, like so many pounds of rice bought by a consumer . . . Of course, one could ask people to state their expectations. That, however, would be simply their guess about how they would act or would have acted in a situation that was not yet at hand or had already passed. Such expectations are no more directly observable than their own preferences and are subject to the same whimsy, arbitrariness, and adjustment to subtle changes in background conditions.

Hoover suggests that expectations fall into the same category as preferences. In revealed-preference theory, a consumer's preference is reconstructed from her behaviour (from her "revealed" preference). Statements about what she thinks she prefers are to be dismissed as neither verifiable nor trustworthy. Similarly, expectation variables cannot be measured because a subject's statement about what she expects can neither be verified nor trusted.

Assume for the sake of argument that a subject's statement about what she expects can be verified or trusted. Manski (2018: section VII) lists a number of problems that would obtain even if that assumption were true: the expectations of firms are difficult to measure; the practice of rounding complicates the interpretation of responses; probabilistic expectations may be ambiguous (to the extent that they express, for instance, ignorance or uncertainty); respondents sometimes confound beliefs and preferences etc. Manski (2018: 423) believes that these problems can be solved. But even if they could be solved, there would be the further problem that survey studies cannot keep track of the way, in which agents form (i.e. revise or update) individual or aggregate expectations. Agents seem to revise or update their expectations more often than they are surveyed. We accordingly need to understand the way, in which they revise or update their expectations. And survey studies cannot keep track of the way, in which they form (i.e. revise or update) individual or aggregate expectations.

One could argue that the way, in which agents form individual or aggregate expectations, can be investigated in laboratory experiments. But laboratory experiments are unlikely to teach us anything about the formation of individual or aggregate expectations. They are unlikely to teach us anything about the formation of individual expectations because in real life, the formation of individual expectations requires that expectations be revised and updated in light of information that (like government announcements, media reports, personal observations etc.) is often generated or collected in obscure ways, and because these obscure ways are difficult (or impossible) to mimic in laboratory experiments. These experiments are also unlikely to teach us anything about the formation of

https://doi.org/10.1017/S0266267125000021 Published online by Cambridge University Press

aggregate expectations because the dynamics of the formation of aggregate expectations in small groups and economic systems have little in common.

Expectations, therefore, cannot be measured. There is currently no way to tell whether the expectations that agents form with respect to the behaviour of macroeconomic variables (or indeed any variable) should be modelled as "rational", "adaptive", or in some alternative way. Expectations are rational if whatever information is available is completely exploited (Muth 1961; Lucas and Prescott 1971). Expectations are adaptive if agents correct past forecasting mistakes when forming expectations.[8]

Critics of rational expectations argue that agents rarely exploit the available information completely, or that they tend to make systematic mistakes when exploiting that information. But proponents of rational expectations respond that systematic mistakes cancel out each other at the aggregate level, and that expectations are correct on average. The critics point out that the response remains unsupported by experimental studies (Anufriev and Hommes 2012) or observational evidence. But the point does not imply that agents never form rational expectations. One may assume, for instance, that in the face of competition, firms have a strong incentive to exploit the available information completely.

In addition to the "adaptive" expectations of consumption-goods producing firms, Dosi *et al.* (2020, section III.C) consider three alternative heuristics that likewise qualify as adaptive. Dosi *et al.* (2020, section V) conduct a number of simulated experiments to show that the stability of macroeconomic system dynamics increases with the degree to which the heuristics of forming demand expectations is "naïve". They also state, however, that their model is able to reproduce the same number of stylized facts under all four heuristics (cf. Dosi *et al.* 2020: 1499). They do not show, that is, how individual agents form expectations as a matter of fact.

Why does the non-measurability of expectations pose a problem for causal inference? It poses a problem because in order to provide conclusive evidence in support of a relation of causal dependence between X and Y, one needs to rule out that there is a confounder: a third variable (or set of variables) Z, on which both X and Y causally depend. In order to rule out that there is such a variable (or set of variables) Z, one needs to control for Z, i.e. include Z in the set of preselected variables and observe that Z remains (largely) unchanged, while X and Y change.

Expectations are often believed to act as confounders in just this sense. Of expectations, however, we cannot know whether they can be controlled for because they cannot be measured. Consider, for instance, the claim that inflation causally depends on aggregate demand. In order to provide conclusive evidence in support of that claim, one would need to rule out that inflation and aggregate demand causally depend on inflation expectations. That they causally depend on inflation expectations is, however, exactly what one would expect when understanding the new Keynesian IS curve and the new Phillips curve of the canonical DSGE model as expressing causal hypotheses. Thus, in order to rule out that inflation and aggregate demand causally depend on inflation expectations, one would need to control for

[8]Friedman (1957) and Phelps (1967) were among the first to apply the conception of adaptive expectations when modeling consumption and the dynamics of inflation and unemployment, respectively.

https://doi.org/10.1017/S0266267125000021 Published online by Cambridge University Press

inflation expectations. The problem is that inflation expectations cannot be measured, and that we cannot tell whether we are able to control for them. In other words: the empirical evidence that we can provide in support of the claim that inflation causally depends on aggregate demand remains inconclusive.

One might think that the case of expectations is an exotic one, and that the non-measurability of expectations does not get in the way of causal inference in macroeconomics in general. But (considered generally and in its most uncontroversial form) the famous Lucas critique says that each of the equations of models representing relations between macroeconomic aggregates needs to be "derived from decision rules … of agents in the economy", that "some view of the behavior of the future values of variables of concern to them …, in conjunction with other factors, determines their optimum decision rules", and that the assumption that this view remains invariant under alternative policy rules is an "extreme assumption" (Lucas 1976: 25).

The Lucas critique says, in other words, that variables Z change whenever policy interventions target X to influence Y, where Z represents an expectational aggregate like inflation or aggregate demand expectations. In the community of agent-based macroeconomists, you sometimes hear voices saying that the empirical relevance of the Lucas critique is questionable in light of the results of empirical procedures of causal inference.[9] But the point is that if Z cannot be measured, the Lucas critique denies that there are reliable empirical procedures that can be employed to provide conclusive empirical evidence in support of causal relations between X and Y, when X and Y stand for macroeconomic aggregates. Thus, the non-measurability of expectations seems to get in the way of causal inference in macroeconomics in general.

5. Sensitive Dependence on Initial Conditions

The second problem arises if the macroeconomy is a chaotic system. I'm aware that a variety of definitions have been proposed for "chaos", and that many of these definitions account for the appearance of randomness in terms of predictability or information theory (cf. e.g. Frigg 2006; Werndl 2009). But fortunately, I don't need to adopt any precise definition of "chaos". I will instead rely on a necessary condition of any notion of chaos: on the condition of sensitive dependence on initial conditions.

The initial conditions of a system can be characterized as a vector of values that the variables of the system attain "initially", i.e. before the interactions between the system components denoted by the variables are iterated. Let $\boldsymbol{S}_0$ be this vector of initial values, and let $\boldsymbol{S}_t$ be the vector of values that the variables of the system attain after t iterations. Then for initial conditions $\boldsymbol{S}_0$, an arbitrary (positive) number $\delta > 0$ and slightly modified initial conditions $\boldsymbol{T}_0$, $|\boldsymbol{S}_0 - \boldsymbol{T}_0| < \delta$ will evolve such that $|\boldsymbol{S}_t - \boldsymbol{T}_t| \approx |\boldsymbol{S}_0 - \boldsymbol{T}_0|e^{\lambda t}$, where λ is the (largest) Lyapunov exponent of the system, which measures the average rate of divergences between neighbouring trajectories that depart from values close to those in $\boldsymbol{S}_0$. We say that a system is (not) sensitive to

[9]Consider, for instance, Delli Gatti *et al.* (2010): "the empirical relevance of the Lucas critique … is largely questionable" (in light of the results of empirical procedures of causal inference).

https://doi.org/10.1017/S0266267125000021 Published online by Cambridge University Press

initial conditions if $\lambda > (<)\ 0$, i.e. if neighbouring trajectories (don't) diverge exponentially. We also say, of course, that the greater (smaller) λ, the more (less) the system is sensitive to initial conditions.

Dosi *et al.* (2013, 2015, 2017) point out on several occasions that they think of the target system of the K+S model as a complex system. There is no universally agreed upon definition of 'complex system', and there is unlikely to be any (single) such definition. But chaos is among the conditions that are regularly proposed as conditions for the complexity of a system. Ask a physicist or engineer what a complex system is, and they are likely to respond that a complex system is chaotic. If chaos is a necessary condition for the complexity of a system, and if the target system of the K+S model is a complex system, then the target system of the K+S model will be chaotic and sensitive to initial conditions, for that matter.

Things are not that easy, however. Ladyman and Wiesner (2020: 78) argue that chaos is not a necessary condition for complexity: that there are complex systems that are not chaotic. Whether a system is chaotic (or sensitive to initial conditions) is, more importantly, an empirical matter. We can use time series for the variables of the system to estimate the Lyapunov exponent: if the estimator is positive, the system will be sensitive to initial conditions; if it is negative, the system will be *in*sensitive to initial conditions. The definition of the Lyapunov exponent involves the infinite time limit (such that, strictly speaking, does not characterize exponential growth for any finite time limit). But an asymptotic distribution of a nonparametric estimator of the Lyapunov exponent can be derived, and that estimator can be used to measure sensitivity to initial conditions (cf. Shintani and Linton 2003). The estimator has been used to reject the null of chaos (or sensitivity to initial conditions) for low-dimensional macroeconomic systems (more specifically, for time series of output in industrialized nations). But the target system of the K+S model (or indeed any macroeconomic ACE model) qualifies as high-dimensional (as including many variables), and the problem with testing for high-dimensional chaos is that one would need extremely long time series to distinguish chaos from randomness. Up to now, these time series are not available for any of the variables included in macroeconomic agent-based models. We accordingly don't know (yet) whether the target system of the K+S model is sensitive to initial conditions.

Why is it important whether the target system of the K+S model is sensitive to initial conditions? Because causal inference will be a lot harder if it is. If the target system of the K+S model is sensitive to initial conditions, then all potential confounding (sets of) variables Z that have been omitted from the model will have to retain exactly the same value z over the course of history that the iterated model is supposed to target. If any of these values deviates from z, the sensitive dependence of the target system on initial conditions will imply that the relation of causal dependence that obtains between X and Y according to the model might fail to have a counterpart in reality.

Consider the relation of causal dependence that obtains between the labour demanded by capital-goods producing firm i at t ($L^D_{i,t}$) and the quantity (machines) produced by that firm at t ($Q_{i,t}$) according to the K+S model. Critics of the K+S model might argue that the model should include a variable Z denoting the demand for the machines that firm i produced at $t-1$ because Z acts as a potential confounder: because both the labour demanded by firm i ($L^D_{i,t}$) and the quantities

https://doi.org/10.1017/S0266267125000021 Published online by Cambridge University Press

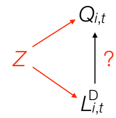

Figure 8. Does the demand for machines act as a confounder?

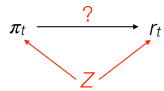

Figure 9. Do inflation expectations act as a confounder?

produced by that firm ($Q_{i,t}$) causally depend on the demand for the machines that firm i produced at t−1 (Z) (Figure 8).

As long as the target system of the K+S model is *in*sensitive to initial conditions, defenders of the K+S model can point out that the demand for the machines produced at t−1 (Z) is negligible: that it can be neglected because the values of Z remain roughly constant over all iterations, while the values demanded labour and produced quantity can be observed to change. But if the target system of the K+S model is *sensitive* to initial condition, then the defenders of the K+S model can no longer claim that Z is negligible: then any slight deviations of the values of Z from z can mean that Z acts as a confounder, and that the relation of causal dependence that obtains between demanded labour and produced quantity according to the K+S model fails to have a counterpart in reality.

Similar considerations apply in the case of the relation of causal dependence that obtains between the rate of inflation (π_t) and the federal funds rate t (r_t) according to the K+S model if the confounding variable Z stands for inflation expectations. Critics of the K+S model might argue that the model should include a variable Z denoting inflation expectations because Z acts as a potential confounder to π_t and r_t (because both π_t and r_t causally depend on Z) (Figure 9).

That Z acts as a confounder to π_t and r_t follows, again, from the canonical DSGE model if the new Philipps curve and the monetary policy rule included in that model are interpreted as expressing causal hypotheses. As long as the target system of the K+S model is *in*sensitive to initial conditions, defenders of the K+S model can point out that inflation expectations are negligible: that Z can be neglected because the values of Z remain the same over all iterations, while the values of π_t and r_t can be observed to change. But if the target system of the K+S model is *sensitive* to initial condition, then the defenders of the K+S model can no longer claim that Z is negligible: then any slight divergences of the values of Z from z can mean that Z acts as a confounder to π_t and r_t, and that the relation of causal dependence that obtains between π_t and r_t according to the K+S model fails to have a counterpart in reality.

Note that I'm not saying that the target system of the K+S model *is* sensitive to initial conditions, or that causal inference is *impossible* if the target system is sensitive to initial conditions. I'm saying that causal inference is harder if the target system of the K+S model is sensitive to initial conditions: that it is the harder, the greater the Lyapunov exponent (λ). The greater the Lyapunov exponent (λ), and the greater the number of iterations (t), the less the values of Z may diverge from a constant value z, and the more potential confounders need to be included in the set

https://doi.org/10.1017/S0266267125000021 Published online by Cambridge University Press

of preselected variables if causal inference is to be successful: if causal inference is supposed to demonstrate that the relation of causal dependence between X and Y does not only obtain in the model.

I would accordingly describe the second problem of causal inference in agent-based macroeconomics as follows: We first need to decide whether the target system of the K+S model is sensitive to initial conditions. If it is, we need to specify the conditions, under which we can say that the Lyapunov exponent and number of iterations are sufficiently small to allow for successful causal inference. The problem is not unsolvable, but it currently seems impossible to decide whether the target system of the K+S model is sensitive to initial conditions (remember that we need long time series to decide that question, and that these time series are currently unavailable). It also seems that there is no general solution to the problem of specifying the conditions, under which we can say that the Lyapunov exponent and number of iterations are sufficiently small.

6. Top-down Causation and Causal Exclusion

The third problem is well-known to philosophers of mind. It arises when philosophers advance the causal exclusion argument in favour of epiphenomenalism, which is the position that mental events supervene on (or are constituted by) physical events and do not cause any physical events (cf. Kim 1989). The causal graph that corresponds to the argument is shown in Figure 10.

Here, P_1 and P_2 are physical events, while M is a mental event, which supervenes on (or is constituted by) P_1. Thus, the fat arrow denotes that P_1 constitutes M (or that M supervenes on P_1). In the original version of the argument, the thin arrows stand for relations of causal sufficiency. And the conclusion of the original version states that M cannot be a cause of P_2 (that there cannot be any top-down causation from M to P_2) if P_1 is a sufficient cause of P_2 because there would be nothing left for M to contribute to the occurrence of P_2.

In a more recent variant of the causal exclusion argument, the thin arrows stand for relations of causal dependence in the sense of Woodward's interventionist definition of causal dependence (cf. section 3 above). In its interventionist version and as expressed by Baumgartner (2010: section III), the causal exclusion argument runs as follows: If P_2 causally depends on P_1, and if M supervenes on P_1, then P_2 cannot causally depend on M because in order for P_2 to causally depend on M, there would have to be a possible intervention on M that changes P_2 while P_1 is held fixed, and because such an intervention is impossible: it is impossible to change M without changing P_1, and vice versa (cf. Baumgartner 2010: section III).

Woodward (2015: sections 6–9) responds that in order for P_2 to causally depend on M, there does not need to be any possible intervention on M that changes P_2 while P_1 is held fixed. He in fact agrees that there cannot be such an intervention. He argues that in order to test whether P_2 causally depends on M, one will have to manipulate an intervention variable that type-level causes both M and P_1, and he expresses the causal dependence of both M and P_1 on the intervention variable by a curly bracket (cf. Woodward 2015: 331) (Figure 11).

https://doi.org/10.1017/S0266267125000021 Published online by Cambridge University Press

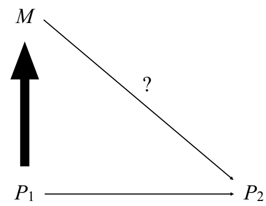

Figure 10. Constitution and causation in the causal exclusion argument.

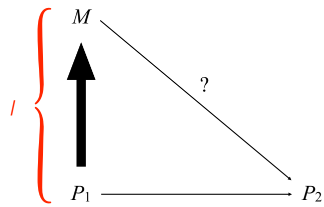

Figure 11. An intervention variable as common type-level cause.

An ensuing change in P_2 would not be attributable to the change in M unambiguously. But an ensuing change in P_2 would at least allow for the *possibility* of there being a process of top-down causation operating from M to P_2.

All sides agree that a process of top-down causation operating from M to P_2 will be impossible if M causally depends on P_1. If M causally depends on P_1, P_1 will act as a confounder and stand in need of control in an experiment carried out to check whether P_2 causally depends on M. Woodward (2015: 316–317) argues that M cannot be regarded as causally depending on P_1 because P_1 and M fail to satisfy the "independent fixability" assumption: because P_1 cannot be manipulated without manipulating M, and vice versa. But Gebharter (2017: section 3) points out that M causally depends on P_1 in the sense that in a causal Bayes net, there will be an arrow departing from P_1 and directed into M (that M will be probabilistically independent of its non-descendants given P_1).

Stern and Eva (2021: section 3) respond that the arrows in a causal Bayes net should be interpreted as representing "e-parenthood": X is an e-parent of Y and Y an e-child of X if either (i) Y causally depends on X, or (ii) Y asymmetrically supervenes on X. If the arrows are interpreted as representing "e-parenthood", Woodward's conclusion, according to which a process of top-down causation operating from M to P_2 is at least possible, is re-established. Stern and Eva (2021: section 6) also point out, however, that in the special sciences, there are cases in which sets of variables violate a condition they refer to as "difference maker (DM) sufficiency". The condition requires essentially that an upper-level variable X screen off a lower-level variable Y from a lower-level variable L that constitutes X: $P(Y|X \wedge L) = P(Y|X)$.[10]

[10] Woodward (2020) mentions a similar condition in a recent contribution to the debate when requiring that an intervention on an upper-level variable U have a uniform effect on Y for all lower-level realizations of the value $U = u$.

As a well-known example of a set violating DM sufficiency, Stern and Eva quote $\{TC, D\}$, where TC and D stand for total cholesterol and heart disease, respectively, and where TC is the sum of high-density cholesterol (HDC) and low-density cholesterol (LDC). $\{TC, D\}$ violates DM sufficiency because TC does not screen off HDC and LDC from D: it matters a great deal to D whether or not LDC is large because LDC influences D negatively, while HDC influences D positively.

In macroeconomics, there are, regrettably, many cases, in which DM sufficiency is violated. Consider, for instance, the set of variables standing for GDP at t and the demand expectations that firm j forms at $t+1$: $\{GDP_t, D^e_{j,t+1}\}$. This set violates DM sufficiency because, arguably, GDP_t does not screen off the quantity produced by firm j at t ($Q_{j,t}$) from the demand expectations formed at $t+1$ ($D^e_{j,t+1}$): it matters a great deal whether or not $Q_{j,t}$ is large because $Q_{j,t}$ is likely to influence demand expectations formed at $t+1$ more strongly than the quantities that any of the other capital- or consumption-goods producing firms produce in t to constitute GDP_t (Figure 12).

More macroeconomic examples of sets violating DM sufficiency are easy to find. Consider e.g. the set $\{\pi_t, P^e_{m,t+1}\}$, where π_t is inflation in t and $P^e_{m,t+1}$ the price that agents expect to pay for commodity m (e.g. real estate) in $t+1$. This set violates DM sufficiency because π_t does not screen off $P^e_{m,t+1}$ from $P_{m,t}$, where $P_{m,t}$ is the price that agents pay for the same commodity in t: it matters a great deal to $P^e_{m,t+1}$ whether or not $P_{m,t}$ is large because $P_{m,t}$ is likely to influence $P^e_{m,t+1}$ more strongly than most of the other prices that form the supervenience base of inflation in t.

When DM sufficiency is violated, manipulations of X will be too "fat-handed" to allow for the conclusion that Y causally depends on X. Manipulations of TC will be too fat-handed to allow for the conclusion that D causally depends on TC. Similarly, manipulations of GDP_t (or π_t) will be too fat-handed to allow for the conclusion that there is a process of top-down causation operating from GDP_t (or π_t) to $D^e_{j,t+1}$ (or $P^e_{m,t+1}$). For causal search algorithms this means that they will result in directed graphs that fail to contain arrows departing from upper-level variables and directed into lower-level ones. For the LiNGAM that Guerini and Moneta (2017) suggest can be employed to discover the relations of causal dependence obtaining in the target system of the K+S model, this means, in particular, that it will result in a directed graph that fails to contain an arrow departing from GDP_t and directed into $D^e_{j,t+1}$.

7. Conclusion

Causal inference is an important branch of macroeconomic research, and macroeconomic ACE models form the latest generation of macroeconomic models that are meant to represent the relations of causal dependence that macroeconomists believe obtain in the economy. Guerini and Moneta (2017) have developed a sophisticated method of providing empirical evidence in support of these relations: they employ a causal search algorithm (the LiNGAM) to identify the parameters of a SVAR without imposing any theoretical restrictions, and they use a similarity measure to assess the extent to which the relations of causal dependence that generate artificial data coincide with the relations of causal dependence that generate real-world data.

https://doi.org/10.1017/S0266267125000021 Published online by Cambridge University Press

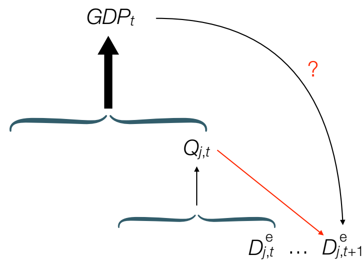

Figure 12. A violation of difference-maker sufficiency in the K+S model.

The inference method that Guerini and Moneta (2017) propose is meant to provide empirical evidence in support of *all* relations of causal dependence that obtain in the target system of macroeconomic ACE models. But so far, they have applied their method only to the macro-part of the K+S model. And the paper has presented three problems that get in the way of an application to the K+S model (or any macroeconomic ACE model) as a whole. By way of conclusion, I will summarize these problems and briefly discuss the steps that I think will have to be taken to bring about an empirical solution to these problems. I will also emphasize that these problems are likely to impede causal inference in macroeconomics more generally, not only in agent-based macroeconomics or when Guerini and Moneta's (2017) inference method is the method of choice.

The first problem is that there is at least one class of hidden variables that cannot be included in the model because they cannot be measured: variables standing for the expectations that agents form about the behaviour of all kinds of micro- or macroeconomic variables that matter to them. The example of inflation expectations acting as a potential confounder to the relation of causal dependence between inflation and aggregate demand shows that this problem does not only arise in agent-based macroeconomics, but in macroeconomics more generally (and macroeconomic DSGE modelling in particular). The problem is solvable in principle. It is not impossible that at one point, we will be able to measure expectations. If we will, we will be able to decide whether expectations can be controlled. And if they can, we will be able to find out whether they truly act as confounders to the relations of causal dependence that we are interested in. I argued that currently, we cannot measure expectations by using survey methods or lab experiments. But maybe survey methods or lab experiments can be improved in such a way that at one point, we can successfully use these methods to measure expectations. Or perhaps we can develop alternative methods: conduct field experiments, collect survey data with a panel structure, or follow individuals over time through their digital footprints. Thus, the first problem seems to be solvable in principle.

The second problem is that of deciding whether the target system of the K+S model is sensitive to initial conditions, and (if it is) that of specifying the conditions under which we can say that the Lyapunov exponent and number of iterations are sufficiently small to allow for successful causal inference. This problem arises in macroeconomics more generally because the target system of the K+S model is the macroeconomy, and because the macroeconomy is the target system of any

https://doi.org/10.1017/S0266267125000021 Published online by Cambridge University Press

macroeconomic model. This problem is also solvable in principle: once the time series available to us are sufficiently long, we can decide whether the target system of the K+S model is sensitive to initial conditions; and once we decide that it is sensitive to initial conditions, we can specify the conditions under which the Lyapunov exponent and number of iterations are sufficiently small to allow for successful causal inference. These conditions are unlikely to hold in general, but perhaps they can be specified on a case-by-case basis.

The third problem is that the target system of macroeconomic ACE models (or indeed macroeconomic models in general) is believed to contain relations of top-down causation (between aggregate quantities and the expectations that individuals of firms form with respect to the behaviour of the micro- or macroeconomic variables that matter to them), that these relations cannot be shown to obtain if the condition of DM sufficiency is violated, and that this condition is likely to be violated in macroeconomics. The example of the causal dependence between inflation at t and the price that agents expect to pay for a commodity at $t+1$ shows that this problem arises in macroeconomics more generally. A solution to this problem requires that macroeconomists understand how agents form individual expectations: that they be able to disentangle the type-level causes of individual expectations (macroeconomic aggregates and microeconomic quantities), and to quantify the degree, to which they are causally relevant. A solution to this problem presupposes a solution to the first problem: macroeconomists won't be able to understand the formation of individual expectations, unless they study the formation of individual expectations empirically, and the empirical study of the formation of individual expectations requires that expectations be measured.

Acknowledgements. The paper is a largely unchanged version of a paper I presented at a workshop on model evaluation and causal search in Pisa in September 2022. I am grateful to Alessio Moneta for written comments on an earlier draft and to the audience for interesting questions and a good discussion. I am also grateful for the comments of two anonymous referees. I gratefully acknowledge financial support from the German Research Foundation (grant: HE 6845/3-1).

References

Anufriev M. and C.H. Hommes 2012. Evolutionary selection of individual expectations and aggregate outcomes. *American Economic Journal: Microeconomics* **4**, 35–64.

Baumgartner M. 2010. Interventionism and Epiphenomenalism. *Canadian Journal of Philosophy* **40**, 359–383.

Cooley T.F. and M. Dwyer 1998. Business cycle analysis without much theory. *Journal of Econometrics* **83**, 57–88.

Delli Gatti D., E. Gaffeo, M. Gallegati, G. Giulioni and A. Palestrini 2008. *Emergent Macroeconomics*. Heidelberg: Springer.

Delli Gatti D., E. Gaffeo and M. Gallegati 2010. Complex agent-based macroeconomics: a manifesto. *Journal of Economic Interaction and Coordination* **5**, 111–135.

Dosi G., G. Fagiolo and A. Roventini 2008. The microfoundations of business cycles. *Journal of Evolutionary Economics* **18**, 413–432.

Dosi G., G. Fagiolo and A. Roventini 2010. Schumpeter meeting Keynes: a policy-friendly model of endogenous growth and business cycles. *Journal of Economic Dynamics and Control* **34**, 1748–1767.

Dosi G., G. Fagiolo and M. Napoletano 2013. Income distribution, credit and fiscal policies in an agent-based Keynesian model. *Journal of Economic Dynamics and Control* **37**, 1598–1625.

Dosi G., G. Fagiolo, M. Napoletano and A. Roventini 2015. Fiscal and monetary policies in complex evolving economies. *Journal of Economic Dynamics and Control* **52**, 166–189.

Dosi G., M. Napoletano, A. Roventini and T. Treibich 2017. Micro and macro policies in Keynes+Schumpeter evolutionary models. *Journal of Evolutionary Economics* **27**, 63–90.

Dosi G., M. Napoletano, A. Roventini, J.A. Stieglitz and T. Treibich 2020. Rational heuristics? Expectations and behaviors in evolving economies with heterogeneous interacting agents. *Economic Inquiry* **58**, 1487–1516.

Friedman, M. 1957. *A Theory of the Consumption Function*. Princeton, NJ: PUP.

Frigg R. 2006. Chaos and randomness: an equivalence proof of a generalized version of the Shannon entropy and the Kolmogorov–Sinai entropy for Hamiltonian dynamical systems. *Chaos, Solitons, and Fractals* **28**, 26–31.

Gebharter A. 2017. Causal exclusion and causal Bayes nets. *Philosophy and Phenomenological Research* **XCV**, 353–375.

Gennaioli N. et al. 2016. Expectations and Investment. *NBER Macroeconomics Annual* 2015, eds. M. Eichenbaum and J. Parker. Chicago: University of Chicago Press.

Glymour C. et al. 2019. Review of causal discovery methods based on graphical models. *Frontiers in Genetics* 10, Art. 525.

Guerini M. and A. Moneta 2017. A method for agent-based models validation. *Journal of Economic Dynamics and Control* **82**, 125–141.

Haavelmo T. 1944/1995. The probability approach in econometrics. In *The Foundations of Econometric Analysis*, eds. D.F. Hendry and M. Morgan, 477–490. Cambridge: Cambridge University Press.

Hitchcock C. and J. Woodward 2003. Explanatory generalizations, Part II: Plumbing explanatory depth. *Nous* **37**, 181–199.

Hoover K.D. 2001. *Causality in Macroeconomics*. Cambridge: Cambridge University Press.

Kim J. 1989. Mechanism, purpose, and explanatory exclusion. *Philosophical Perspectives* **3**, 77–108.

Koopmans T.C. (ed.) 1950. When is an equation system complete for statistical purposes? In *Statistical Inference in Dynamic Economic Models (Cowles Commission Monograph No. 10)*, 393–409. New York: Wiley & Sons.

Kydland F.E. and E.C. Prescott 1982. Time to build and aggregate fluctuations. *Econometrica* **50**, 1345–1370.

Ladyman J. and K. Wiesner 2020. *What Is a Complex System?* New Haven: Yale University Press.

Lucas R.E. 1976. Econometric policy evaluation: a critique. *Carnegie-Rochester Conference Series on Public Policy* **1**, 19–46.

Lucas R.E. and E.C. Prescott 1971. Investment under uncertainty. *Econometrica* **39**, 659.

Manski C.F. 2018. Survey measurement of probabilistic macroeconomic expectations. *NBER Macroeconomics Annual* **32**, 411–471.

Muth J.F. 1961. Rational expectations and the theory of price movements. *Econometrica* **29**, 315–353.

Phelps E.S. 1967. PCs, expectations of inflation, and optimal unemployment over time. *Economica* **34**, 254–281.

Shintani M. and O. Linton 2003. Non-parametric neural network estimation of Lyapunov exponents and a direct test for chaos. *Journal of Econometrics* **120**, 1–33.

Sims C.A. 1980. Macroeconomics and reality. *Econometrica* **48**, 1–48.

Stern R. and B. Eva 2023. Antireductionist interventionism. *British Journal for the Philosophy of Science* **74**, 241–267.

Werndl C. 2009. What are the new implications of chaos for unpredictability? *British Journal for the Philosophy of Science* **60**, 195–220.

Woodward J. 2003. *Making Things Happen: A Causal Theory of Explanation*. Oxford: Oxford University Press.

Woodward J. 2015. Interventionism and causal exclusion. *Philosophy and Phenomenological Research* **91**, 303–347.

Woodward J. 2020. Causal complexity, conditional independence, and downward causation. *Philosophy of Science* **87**, 857–867.

Tobias Henschen is a principal investigator in a research project that is concerned with the philosophical foundations of complexity economics, funded by the German Research Foundation, and hosted by the Philosophy Department of the University of Cologne. Previously, he held positions at University College Freiburg and at the Philosophy Departments of the Universities of Konstanz and Heidelberg. He obtained degrees in philosophy and economics from the Universities of Toulouse 1 and Heidelberg and a venia *legendi* (*Habilitation*) for philosophy at the University of Konstanz. His main research interests include general philosophy of science, philosophy of economics, and history of philosophy. URL: https://www.tobiashenschen.net

Cite this article: Henschen T. Some problems of causal inference in agent-based macroeconomics. *Economics and Philosophy*. https://doi.org/10.1017/S0266267125000021