Energy Conversion and Management: X 32 (2026) 102317

Contents lists available at ScienceDirect

Energy Conversion and Management: X

journal homepage: www.sciencedirect.com/journal/energy-conversion-and-management-x

Convergence, certified solutions, and equilibrium enumeration for scalar-linked top-down and bottom-up energy-economy models

Sergio Bobka, Julia Gershenzon, Stefan Vögele, Jochen Linßen, Andrew G. Ross *

Forschungszentrum Jülich GmbH, Institute of Climate and Energy Systems, Jülich Systems Analysis, Wilhelm-Johnen-Strasse, 52425 Jülich, Germany

ARTICLE INFO

Dataset link: https://github.com/ag-ross/TDBU

JEL classification:
C63
C68
Q40
C61

Keywords:
Model linking
Top-down and bottom-up
Computable general equilibrium
Integrated assessment
Parametric linear programming
Multiple equilibria

ABSTRACT

Linked energy-economy models can iterate between a top-down model of the economy and a bottom-up model of the energy system. This process continues, with each model exchanging relevant quantities, until their results converge. Although a recent coordination algorithm lowers the cost of each round by predicting bottom-up solutions from stored bases, the alternation remains a fixed-point iteration. It converges only where the round map is locally contracting, its stopping rule bounds no distance to a fixed point, and with several equilibria it reports whichever one the starting point selects. This paper addresses all three shortcomings for linkages that exchange a single scalar in each direction. The linked equilibrium becomes the root of a scalar function, located by bracketing with an exact derivative recovered from the stored basis inverse, and every candidate carries either a certificate that a true equilibrium exists nearby and is locally unique, or an explicit failure report. Extending the pointwise basis test to its closed-form validity interval makes the bottom-up response exactly piecewise rational, so all equilibria in a stated range along the tracked top-down branch can be enumerated and classified. The scalar solver reproduces the published equilibria, each unique in that range, with half the top-down solves and slightly more bottom-up solves. A parametric root finder removes almost all bottom-up work. The gain at which the corner equilibrium is born has a closed form built from the validity intervals. An amplified demand link illustrates the structure beyond that threshold, including a repelling equilibrium that no damped alternation can reach.

1. Introduction

Energy system models with technological detail are routinely linked to macroeconomic models, a combination the hybrid-modelling literature treats as a subject in its own right [6]. In a hard link the two models are combined into a single system of equations and solved simultaneously, so the exchanged quantities are consistent by construction and a solution is a point satisfying the combined system rather than an iterate stopped on a tolerance [e.g. 2,7]. A soft link instead keeps the models as separate programs, each solved on its own and coupled by exchanging data between the solves. The exchange may run in one direction, with one model taking inputs from the other and nothing passed back, or in both, as here, where the models are solved in alternation and information passes each way until the exchanged quantities settle [2,8–11]. Iteration between separate solves, not the direction of exchange, is what sets a soft link apart from the single combined solve of a hard link. In practice, though, hard-linking is rarely available for the models of interest. A technology-rich energy system model such as TIMES[1] is large, developed independently of the economic model and often accessible only through a solver rather than as an explicit set of equations, so merging the two into one system is impractical and the coupling proceeds by soft-linking instead [12].

Pisciella et al. [3] (PvBT hereafter) exploit the fact that repeated solves of the bottom-up (BU) linear programme (LP) account for the bulk of the computational cost in linked models of this kind, the top-down (TD) block amounting to only a few dozen market-clearing and zero-profit conditions [2,7,9].[2] They propose a coordination algorithm

* Corresponding author.
E-mail address: a.ross@fz-juelich.de (A.G. Ross).

[1] TIMES (The Integrated MARKAL-EFOM System) is an energy system optimisation framework [1] that models the energy system as a linear programme over technologies, commodities and time periods. The instance used in the numerical experiments of this paper, taken from [2,3], follows TIMES structural conventions but is far smaller than a production deployment, with seven processes and 144 structural columns against the thousands of technologies in a national-scale model.

[2] The zero-profit condition of competitive computable general equilibrium (CGE) models states that unit cost equals price wherever output is positive [4]. The condition presupposes constant returns and price-taking. Both hold in the TD block of Section 2, where factor demands (4) and intermediate use are linear in output and the second equation of (6) carries no markup. Because the cost side of that equation includes the return to capital, accounting profit may still be positive. Dixon and Rimmer [5] give a contrasting case, a global CGE model in which some industries are modelled under small-group monopolistic competition with a sticky number of firms, and blocked or partially blocked entry lets incumbents earn positive pure profits.

https://doi.org/10.1016/j.ecmx.2026.102317
Received 24 July 2026; Received in revised form 7 September 2026; Accepted 12 September 2026
Available online 18 September 2026

that predicts the LP solution from previously stored optimal bases. When the prediction passes a sign test it equals the exact LP optimum and the solve is skipped. Their numerical experiments report sizeable time savings. The present analysis reproduces the algorithmic content of that claim, namely the large reduction in the number of BU solves, in an independent implementation.

Acceleration does not, however, address reliability, and PvBT do not claim it does. Their contribution concerns the cost of each round, and reliability is a question their paper does not pose. The alternation that PvBT accelerate is a fixed-point iteration, and fixed-point iterations converge only to attracting fixed points, stop on heuristics, and select among multiple equilibria by chance of the starting point. PvBT note themselves that oscillating patterns 'cannot be excluded' a priori and that Helgesen and Tomasgard [2] observe such patterns in practice. This paper asks what can be guaranteed, rather than solely accelerated, and answers the question for the linkage class used in the PvBT experiments, up to the stated caveats on grid resolution, degeneracy, and the single-valuedness of the top-down response (see Appendix A for a more detailed overview). The result is a complement to the PvBT acceleration, which it leaves intact, not a replacement for it.

The reliability questions posed here have a history in the linking literature. Wene's [8] early comparison of linking strategies distinguishes formal convergence of the exchanged quantities from the ad hoc stopping practice of linked studies, and Krook-Riekkola et al. [9] document how difficult it is, in a full TIMES-CGE soft link, to decide when convergence has been reached. Andersen et al. [13] restructure the exchanged quantities themselves, linking through energy service demands to improve the stability of the iteration. On the integration side, the complementarity treatment of economic equilibria of Mathiesen [4] underlies the single-system formulations of Böhringer and Rutherford [7,12], which Section 4.5 uses both as a baseline solver and as the certificate evaluated at every candidate solution.[3] Acceleration schemes for fixed-point iterations, Anderson [14] mixing among them, can converge locally even where the damped alternation diverges, but like the alternation they certify no limit point and reveal nothing about further equilibria. The validity intervals at the centre of the present scheme are the one-parameter case of multiparametric linear programming, classical since Gal and Nedoma [15] and algorithmic in modern form in Borrelli et al. [16].

The analysis proceeds in four steps. First, the linked equilibrium is reduced to a root of a scalar function, because the information exchanged between the two models consists of one scalar in each direction. A bracketing method then converges to an equilibrium without any contraction assumption. Second, an exact derivative is available without finite differences, since the stored basis inverse of PvBT doubles as the sensitivity of the BU response and the implicit function theorem supplies the sensitivity of the TD response. Third, every returned solution carries a certificate that a true equilibrium exists nearby and is locally unique, obtained by re-evaluating the residual without predictions and applying a Newton–Kantorovich (NK) test to the integrated equilibrium system (the TD equations together with the LP optimality conditions and the linking equations). Where the certificate does not hold, an explicit failure report replaces it. Fourth, and centrally, the pointwise prediction test of PvBT is extended to a closed-form validity interval for each basis, a specialisation of classical right-hand-side ranging [15] to the linkage. The novelty is the use made of the interval for the linked problem. The BU response then becomes an exactly computable piecewise rational function, and the set of all equilibria in a range can be enumerated and classified. The numerical section shows that this changes the qualitative conclusions available from the same model. At the published calibration the equilibrium found by the GAMS code is the only one in the economically relevant range, a statement established by decomposing that range into exactly known smooth pieces and scanning each piece with its analytic derivative (the remaining failure modes, a root pair closer than the refinement grid inside one piece and the branch assumption on the CGE response, are stated in Appendix A). By contrast, a stressed calibration of the same model possesses three equilibria, of which the alternation can only ever reach two.

All computations are carried out in a dependency-free Python implementation that reproduces the published GAMS results of PvBT before any modification is introduced. Section 2 gives the background, stating the models, the linkage and the PvBT algorithm and formalising the shortcomings the alternation leaves. Section 3 develops the solution scheme, Section 4 applies it to the CGE-TIMES linkage, Section 5 discusses the results and their limits, and Section 6 concludes, with the finer numerical caveats collected in Appendix A and the symbols in Appendix B.

2. Conceptual background

Introducing no new method of its own, this section reviews existing work and the limitations it leaves. It states the abstract framework of the linked problem and the PvBT coordination algorithm, describes the CGE-TIMES instance used as the test case throughout, and sets out the reliability shortcomings of the existing alternation. The scheme that closes those shortcomings is the paper's own contribution and is developed separately in Section 3. Appendix B groups the notation of the whole paper by block.

2.1. General formulation and the coordination algorithm

PvBT consider a TD model in mixed-complementarity (MCP)[4] form and a BU model in LP form,

$$\mathrm{TD}(y): \quad 0 \le F(x,y) \;\perp\; x \ge 0, \tag{1}$$

$$\mathrm{BU}(x): \quad \min_{y \in \mathbb{R}^n_+} \left\{ c^\top y \;\middle|\; Ay = b(x) \right\}, \tag{2}$$

where the TD solution x enters the BU model only through the right-hand side $b(x)$ (here c is the LP cost vector, A the constraint matrix, n the number of bottom-up variables, and $\perp$ in (1) denotes complementarity), and the BU solution y enters the TD model through a fixed

[3] Certification addresses the stopping rule. The standard alternation halts when successive iterates change by less than a tolerance, which establishes only that the iteration stalled. Certification replaces this with a direct check. That is, once a candidate solution is produced, its residual in the full integrated equilibrium system is evaluated independently of the path that led to it. The check assembles whether the candidate satisfies the top-down equilibrium conditions, the bottom-up optimality conditions and the linking equations simultaneously, each to a stated numerical bound. The output is not that the iteration stopped but that a specific point satisfies, or fails to satisfy, the equilibrium system to a quantified tolerance.

[4] The MCP form of (1) is standard in the Walrasian competitive tradition of Mathiesen [4] and Böhringer and Rutherford [7,12], but it is a presentational choice rather than a restriction. Many features that appear to lie beyond it admit complementarity representations within the same tradition, imperfect competition through markup conditions, involuntary unemployment through a wage-setting inequality with unemployment as the complementary slack, and intertemporal models from perfect foresight to recursive myopia through their stationarity conditions. What the present scheme requires is not the complementarity format as such but the properties assumed of the top-down block, namely that its response is a single-valued and differentiable function of the one exchanged scalar, so that the linkage reduces to a scalar root with an available derivative. At the interior equilibria of interest the MCP is in any case solved as the smooth square system (3)–(6), and a non-complementarity solve of the same conditions, as in the models of [e.g. 17–19], carries the argument unchanged where those properties hold. The guarantees lapse where the properties fail, particularly under forward-looking dynamics that exchange a trajectory rather than one scalar, or responses that are not single-valued. This paper addresses only the competitive formulation used by PvBT and makes no claim about others.

functional link. A linked equilibrium is a pair (x^*, y^*) solving both problems simultaneously. The standard procedure alternates between (1) and (2) until the exchanged data change by less than a tolerance. If B is an optimal basis of (2) at $\bar{x}$, the basic part of the solution is $y_B = B^{-1}b(\bar{x})$ with $y_N = 0$. The observation of PvBT is that for a nearby x the same basis often remains optimal, in which case the LP solution is available by a matrix–vector product.

PvBT's Theorem 1 states that the prediction $\hat{y}_B = B^{-1}b(x)$ is optimal if and only if $\hat{y}_B \geq 0$ (dual feasibility is unaffected because c does not depend on x).[5] The coordination algorithm stores the bases of all previous LP solves together with their optimal dual vectors λ_k for (2), selects among them by maximising the dual objective $\lambda_k^{\mathsf{T}} b(x)$, and solves the LP from scratch only when the sign test fails. The cut selection is itself embedded in the TD problem as an MCP block, the 'top-down master', with a scalar variable θ and weights ζ_k representing $\theta = \max_k \lambda_k^{\mathsf{T}} b(x)$. The algorithm applies to any TD-BU linkage of the general form (1)–(2). The specific instance used as the test case throughout this paper is introduced next.

2.2. The CGE-TIMES test instance

The numerical experiments of PvBT link a simplified four-sector CGE model to a simplified TIMES LP, both taken from [2]. The CGE has sectors $S = \{\mathrm{GAS}, \mathrm{ELE}, \mathrm{MAN}, \mathrm{NON}\}$, one household with linear-expenditure-system (LES) preferences, constant elasticity of substitution (CES) value added over capital and labour, and a Leontief intermediate block with technical coefficients io_{ij}. With P_K the return to capital, P_L the wage (fixed as numeraire), P_i the commodity prices, X_i gross output, $\bar{K}, \bar{L}$ the capital and labour endowments, K_i, L_i the factor demands, C_i consumption, and R income, and with α_i the marginal expenditure share and μ_i the subsistence consumption in the LES, γ_i the capital weight and σ_i the substitution elasticity in the CES value added, and a_i the CES scale parameter, the equilibrium conditions are

$$P_i C_i = P_i \mu_i + \alpha_i \big(R - \textstyle\sum_j \mu_j P_j\big), \qquad i \in S, \tag{3}$$

$$K_i = \frac{X_i}{a_i}\Big(\frac{\gamma_i}{P_K}\Big)^{\sigma_i} \Phi_i^{\sigma_i/(1-\sigma_i)}, \qquad L_i = \frac{X_i}{a_i}\Big(\frac{1-\gamma_i}{P_L}\Big)^{\sigma_i} \Phi_i^{\sigma_i/(1-\sigma_i)}, \qquad i \in S, \tag{4}$$

$$\bar{K} = \textstyle\sum_i K_i, \qquad X_i = C_i + \sum_j \mathrm{io}_{ij} X_j, \qquad i \in S, \tag{5}$$

$$R = P_K \bar{K} + P_L \bar{L}, \qquad P_K K_i + P_L L_i + \textstyle\sum_j \mathrm{io}_{ji} P_j X_i = P_i X_i, \qquad i \in S, \tag{6}$$

with $\Phi_i = \gamma_i^{\sigma_i} P_K^{1-\sigma_i} + (1-\gamma_i)^{\sigma_i} P_L^{1-\sigma_i}$, and labour-market clearing implied by Walras' law. At the equilibria of interest all variables are strictly positive, so the MCP (1) reduces to the square system (3)–(6).

The BU model is a twelve-period LP with seven processes (hydro and gas power, gas burners, electric heaters, gas production and two demand processes), 132 constraint rows plus an objective accounting row, and 144 structural columns. It minimises investment, fixed, variable and commodity costs net of salvage values subject to capacity, commodity-balance, production-accounting and investment-bound constraints. Only its demand rows depend on the TD model. The linkage exchanges one scalar in each direction. The TD model passes a demand multiplier

$$D = \frac{X_{\mathrm{GAS}} + X_{\mathrm{ELE}} - X^0_{\mathrm{GAS}} - X^0_{\mathrm{ELE}}}{X^0_{\mathrm{GAS}} + X^0_{\mathrm{ELE}}}, \tag{7}$$

where X^0_{GAS} and X^0_{ELE} are the benchmark (calibration) gross outputs of the gas and electricity sectors, so that D measures the proportional deviation of combined gas and electricity output from its benchmark level and vanishes at the calibration point. This multiplier scales the BU demand projections linearly over the horizon, so that the LP right-hand side is affine, $b(D) = b_0 + D\,\Delta b$ (with $b_0 = b(0)$ and $\Delta b = \mathrm{d}b/\mathrm{d}D$). The BU model passes back the gas-to-electricity technology coefficient

$$\iota = \mathrm{io}_{\mathrm{GAS,ELE}} = \frac{\mathrm{ACT}_{\mathrm{gaspower},\,2020}}{\mathrm{ACT}_{\mathrm{ele\text{-}dem},\,2020}}, \tag{8}$$

the ratio of gas-fired generation to electricity demand activity in the final period (with $\mathrm{ACT}_{p,t}$ the activity level of bottom-up process p in period t, the subscript 2020 labelling the final period of the horizon), which replaces one entry of the Leontief block in (5) and (6). Throughout the paper ι (Greek iota) denotes this scalar, distinct from the full coefficient matrix io_{ij} of (5)–(6). Convergence is declared when D, the LP objective and a derived quantity change by less than 10^{-6} in relative terms between successive rounds. With the model and linkage fully defined, the next subsection examines what the alternation guarantees and what it leaves open.

2.3. Shortcomings of the alternation

Three reliability shortcomings of the alternation as a fixed-point iteration motivate the scheme developed in Section 3, together with a fourth observation on the structure of the augmented TD problem. Define the two half-round maps induced by the linkage,

$$f : \iota \mapsto D \quad \text{(solve the CGE at } \iota\text{, apply (7))}, \qquad g : D \mapsto \iota \quad \text{(solve the LP at } D\text{, apply (8))}, \tag{9}$$

and the round map $\varphi = g \circ f$. The alternation, accelerated or not, is the fixed-point iteration $\iota_{k+1} = \varphi(\iota_k)$, and a linked equilibrium is a fixed point $\iota^* = \varphi(\iota^*)$. Three shortcomings follow from elementary fixed-point theory, and a fourth from the construction of the augmented TD problem.

The first shortcoming is conditional convergence. A standard result from fixed-point theory states that the iteration $\iota_{k+1} = \varphi(\iota_k)$ is locally convergent to a differentiable fixed point ι^* if $|\varphi'(\iota^*)| < 1$ and locally divergent if $|\varphi'(\iota^*)| > 1$. This condition is not checked, and cannot be checked, by the alternation itself. Section 4 measures φ' at the published calibration (about 0.47 at the benchmark point of the $s_K = s_L = 1.1$ case, where s_K and s_L are the multiplicative factors that scale the benchmark capital and labour endowments $\bar{K}, \bar{L}$ across the 7×7 shock grid of PvBT, and between 0.33 and 0.59 at the 49 reported equilibria), which explains why PvBT observe no failures, and exhibits a stressed calibration of the same model with $\varphi'(\iota^*) \approx 13$ at one equilibrium that no damped alternation can then reach. Damping does not help at such a point, since the damped map $(1-\omega)\,\iota + \omega\,\varphi(\iota)$ has slope $1 + \omega(\varphi' - 1) > 1$ at the fixed point for every $\omega \in (0, 1]$ whenever $\varphi' > 1$. This obstruction is specific to the sign of φ'. Where divergence is instead oscillatory, $\varphi' < -1$, any damping factor $\omega < 2/(1-\varphi')$ restores local convergence, so a repelling point of that kind is reachable by a damped alternation, whereas the stressed equilibrium has $\varphi' > 1$ and is not. Anderson-type acceleration [14] can converge locally even at a point of this kind, but it shares the alternation's silence on certification and on the existence of further equilibria, which are the second and third shortcomings below.

The second shortcoming is the stopping rule. The alternation stops once successive rounds differ by less than a tolerance. A small step does not bound the distance to a fixed point unless a contraction modulus is known, and the specific rule used in the GAMS implementation, $100\,|(\mathrm{new} + 10^{-3})/(\mathrm{prev} + 10^{-3}) - 1| < 10^{-4}$ applied to three monitored quantities, also degenerates when a monitored quantity approaches -10^{-3}. Termination therefore certifies that the iteration stalled, not that an equilibrium has been found.

[5] The fixed cost vector is also a scope condition. Linkages that feed the top-down solution back into the bottom-up objective, through marginal costs, fuel prices or a carbon price, make c depend on x and fall outside the class treated here, as does the elastic-demand formulation of MARKAL-ED [20], in which demand response enters the objective itself. Section 5 returns to this boundary.

The third shortcoming is multiplicity. When φ has several fixed points the alternation selects one by chance of its starting point and reports nothing about the others. Repelling fixed points, which separate the basins of attracting ones, are invisible in principle. In an energy-economy linkage such points are the boundaries between qualitatively different technology regimes, so their location is informative rather than incidental. Attracting and repelling are properties of the algorithmic round map φ, not of tâtonnement or any other market adjustment process, and carry no implication about the dynamic stability of the underlying economy. Nor is multiplicity an exotic concern in general equilibrium, as uniqueness fails in well-behaved economies [21], and in applied models it has historically been settled only case by case, by numerical search [22].

In the augmented TD problem of PvBT the variables θ and ζ_k appear in no equation other than the cut block itself. The augmented formulation therefore decouples, with its CGE part coinciding with the plain TD solution and (θ, ζ) encoding $\theta = \max_k \lambda_k^{\mathsf{T}} b(x)$ with ζ supported on the argmax. This is convenient in what follows, because the cut selection can be computed directly, and it also implies that the augmented formulation brings no additional fixed-point structure that might aid convergence.

Two side effects observed in the GAMS code are worth recording. The first is minor. With an empty cut list the constraint $1 = \sum_k \zeta_k$ is infeasible, though no failure arises in practice because the solver's presolve removes the row before it is seen. The second is more substantive. The prediction step assembles the technology coefficient from the basic-variable ordering of the most recent LP solve, even when the active cut ζ selects an older basis B_k. The sign test is unaffected, but the coefficient assembled in this way would be incorrect whenever the selected basis differs from the latest one in its basic-variable ordering.[6] PvBT report that in their runs the active basis is always the most recent one, so their published results are unaffected, yet the oscillating patterns they cannot exclude a priori are exactly the case in which old bases are selected. The port used in Section 4 reproduces this behaviour for replication purposes.

3. Methods

This section develops a certified bracketing and enumeration scheme (CBE), the contribution of this paper, that closes the three shortcomings of Section 2.3 while reusing, rather than replacing, the per-round acceleration of PvBT. It proceeds in four steps. The first reduces the linked equilibrium to a root of a scalar function, turning the fixed-point iteration into a bracketing problem and thereby removing the contraction requirement entirely. The second shows that the exact derivative of that scalar function is available from the stored basis inverse that PvBT already maintain, enabling a Newton-accelerated bracketing method at no additional factorisation cost. The third certifies candidate solutions, replacing the heuristic stopping rule with a Newton–Kantorovich test on the integrated equilibrium system that establishes existence of a true equilibrium nearby and its local uniqueness, or fails explicitly. The fourth and central step extends the pointwise basis test of PvBT to a closed-form validity interval for each stored basis. This makes the BU response an exactly known piecewise rational function, so all equilibria in a stated range, on the tracked top-down branch (see Appendix A), can be enumerated and classified, and the gain at which the corner equilibrium is born has a closed form.

Fig. 1 charts the scheme as implemented. Its shaded panel encloses one round of PvBT's accelerated alternation, which feeds both columns beneath it. Traced with the test instance, the round starts from a trial

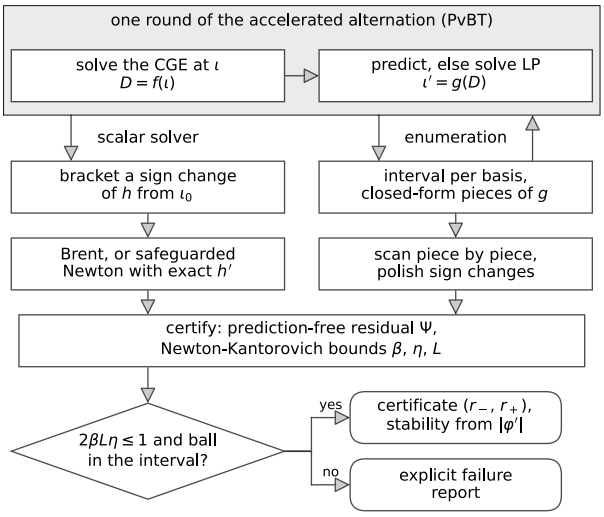

Fig. 1. Certified bracketing and enumeration as implemented. The shaded panel is one round of the accelerated alternation of [3], in which a stored basis predicts the LP solution whenever its sign test passes. Below it, the left column locates one equilibrium per bracket by Brent's method (Section 3.1) or by safeguarded Newton steps on the exact derivative (Section 3.2). The right column extends the sign test to the validity interval (12), whose chained pieces return g to the round in closed form and align the scan that finds every root in the range (Section 3.4). Candidates from both columns pass through the certification of Section 3.3 and leave with a Newton–Kantorovich certificate or an explicit failure report.

value of the technology coefficient ι and passes it to the CGE, whose solution at that value gives the demand multiplier $D = f(\iota)$ of (7). The stored basis that PvBT's cut selection picks at that demand then supplies the LP solution by prediction if its sign test passes, and the LP is solved if it fails.[7] In either case, (8) returns the next coefficient $g(D) = \varphi(\iota)$. Panel (a) of Fig. 2 shows this round map φ at the published calibration. It crosses the diagonal once, at the linked equilibrium, with a slope below one, so the alternation converges from nearby starts. The two columns of Fig. 1 differ in what they return. The scalar solver on the left brackets that crossing, as a sign change of the one-round update h of (10), instead of iterating towards it, and locates one equilibrium per bracket by Brent's method or by safeguarded Newton steps on the exact derivative (Sections 3.1 and 3.2). The enumeration on the right returns every equilibrium in a stated range of ι. It chains the pieces of Fig. 4, on each of which one basis stays optimal, into a cover of the image of that range under f. The return arrow carries these pieces back into the panel as g in closed form. Splitting the range of ι at the preimages of the breakpoints keeps h smooth on every scanned cell (Section 3.4). Candidates from either column then pass through the certification of Section 3.3. Sections 3.1 to 3.4 develop the four steps in turn, and Section 3.5 sets out how the scheme as a whole relates to the coordination algorithm of PvBT.

3.1. Scalar reduction and guaranteed convergence

Because the linkage exchanges one scalar each way, a linked equilibrium is a root of

$$h(\iota) = g\big(f(\iota)\big) - \iota. \tag{10}$$

[6] A further defect of the same family sits in the adjacent line of the GAMS source, where the slack-count offset that anchors this mapping is summed over all cuts with positive weight, so a tie in the cut selection would corrupt the offset itself.

[7] The scalar solver of Section 4 tries the three stored bases with the largest dual objective before solving the LP and keeps the 25 most recent bases, whereas PvBT test one basis and keep all. The port of the coordination algorithm follows PvBT's rule exactly. The bottom-up counts of the Brent and Newton variants in Section 4.4 are obtained under the solver's rule.

Root finding for a scalar function is a solved problem. If h is continuous on $[\underline{\iota}, \bar{\iota}]$ and changes sign, a bracketing method converges to a root, and Brent's method [23] does so superlinearly in practice while retaining the bisection guarantee. No contraction property is required, so repelling fixed points are found as readily as attracting ones. The bracket is obtained by expanding from the benchmark value of ι in steps scaled by $|h|$, which is the natural step length because h is itself the one-round update.

3.2. Derivatives without finite differences

A safeguarded Newton variant accelerates the bracketing method using the exact derivative of h, which is available from objects already held in the PvBT machinery. On the BU side, within an optimal basis the solution is affine in the right-hand side, $y_B(D) = B^{-1}b_0 + D\,B^{-1}\Delta b$, so the activity ratio (8) is rational in D and

$$g'(D) \;=\; \frac{v_1\,(u_2 + v_2 D) - (u_1 + v_1 D)\,v_2}{(u_2 + v_2 D)^2}, \tag{11}$$

where $u = B^{-1}b_0$ and $v = B^{-1}\Delta b$, and (u_1, v_1) and (u_2, v_2) are their components at the basis positions of $\mathrm{ACT}_{\text{gaspower, 2020}}$ and $\mathrm{ACT}_{\text{ele-dem, 2020}}$, both basic at the interior equilibria. The stored basis inverse that PvBT use to predict solutions is thus also an exact derivative oracle, a zeroth-order device upgraded to first order at no additional factorisation cost. On the TD side, with J_{CGE} the analytic Jacobian of the square CGE system (3)–(6) with respect to the 22 CGE unknowns $z = \left(P_K,\, R,\, \{P_i, X_i, K_i, L_i, C_i\}_{i\in S}\right)$ (P_L excluded as numeraire), and s the derivative of the CGE residual with respect to ι, the implicit function theorem gives $\mathrm{d}z/\mathrm{d}\iota = -J_{\mathrm{CGE}}^{-1}s$,[8] in which s has exactly two non-zero entries (one in the market-clearing row for gas, one in the zero-profit row for electricity). Then f' follows from (7) and $h' = g'(f(\iota))\,f'(\iota) - 1$ exactly on each smooth piece. Off-bracket Newton steps are replaced by bisection, so the convergence guarantee of the bracket is retained.

3.3. Certification

Whichever method produces a candidate $\tilde{\iota}$, the implementation re-evaluates $h(\tilde{\iota})$ with the prediction device switched off (one clean CGE solve and one clean LP solve, both included in the work counts of Section 4) and assembles the residual of the integrated equilibrium system, namely the CGE Eqs. (3)–(6), primal and dual feasibility of the LP (the latter comprising both the structural reduced costs and the sign of the inequality-row multipliers), complementary slackness, and the two linking Eqs. (7)–(8). The assembly itself is arithmetic on the two solutions in hand, and verifying a certificate never requires a further model solve. When a stored basis covers the candidate the clean LP solve can also be dispensed with, since the predicted primal solution and the stored duals supply every quantity the residual needs. This is the single-MCP formulation of the linked model in the sense of Böhringer and Rutherford [12], evaluated as a certificate rather than solved. The object certified is therefore a solution of the linked system exactly as PvBT pose it, with the exchange map of Section 2 taken as given rather than defended as the right embedding of the energy system in the economy.

A small residual alone bounds no distance to a true equilibrium. On a fixed optimal LP basis the stacked integrated system $\Psi(w) = 0$ (with Ψ the full integrated equilibrium residual, distinct from the top-down complementarity map F of (1)), in the 23 unknowns $w = (z, \iota)$, is smooth, so a Newton–Kantorovich argument [24] turns the residual into an existence-and-uniqueness certificate at a candidate $\tilde{w}$. Write $J = \partial\Psi/\partial w$ for the 23×23 Jacobian of this stacked system (distinct from the 22×22 CGE Jacobian J_{CGE} above), η for a verified upper bound on the first Newton step length $\|J^{-1}\Psi(\tilde{w})\|_\infty$, β for a verified upper bound on $\|J^{-1}\|_\infty$ obtained from $\|I - CJ\|_\infty < 1$ for the computed inverse C [25] with the floating-point product-error term [26] included, and L for a certified enclosure of the Lipschitz constant of J over a ball about $\tilde{w}$, formed by an interval extension in P_K with the remaining blocks in closed form. When $2\beta L\eta \le 1$ and the ball's demand excursion lies within the basis validity interval (12), a true equilibrium exists within radius $r_- = (1 - \sqrt{1 - 2\beta L\eta})/(\beta L)$ of $\tilde{w}$ and is the unique equilibrium at distance less than $r_+ = (1 + \sqrt{1 - 2\beta L\eta})/(\beta L)$ (uniqueness holds on the open ball, and the quadratic majorant itself has a second zero at exactly r_+). Because uniqueness is asserted on the ball of radius r_+, the enclosure is taken over a ball at least that large rather than over the Newton step, its radius being enlarged and L recomputed until the two are consistent.

The violations quoted for this residual throughout the paper are maxima of absolute values over blocks in their natural units, with no rescaling applied, so a stated bound is most demanding on the rows of largest natural scale. The Newton step $J^{-1}\Psi$, and with it η, is invariant to a rescaling of the equations and is measured in the units of the unknowns, while β and L are taken in the same fixed natural units throughout. The residual, the Jacobian, the Lipschitz enclosure and the interval margins are all evaluated in floating point without directed rounding, so the certificate is rigorous up to that evaluation error, whose reach the reported quantities bound. The stacked Jacobian is well conditioned, with $\mathrm{cond}_\infty(J)$ below 1.5×10^4 at every certified candidate on the 49-case grid and in the stressed calibration of Section 4, and the residual rows are assembled from terms no larger than order 10^2, so a standard forward-error argument perturbs the computed $2\beta L\eta$ by less than 10^{-6} in absolute terms. Since the worst case reported below passes the test at $2\beta L\eta \le 10^{-6}$ against a threshold of one, the verdict survives this perturbation with five orders of magnitude to spare, while the quoted existence radii carry order-of-magnitude accuracy. Outward rounding of the same quantities would remove the qualification entirely at negligible cost for a system of this size, and is a hardening the implementation does not include. The certificate then proves that a true equilibrium exists within a computed radius of the candidate and is locally unique. The test holds at every interior equilibrium reported in Section 4, with worst-case $2\beta L\eta$ below 10^{-6} and existence radius below 2×10^{-10} over the 49-case grid. The corner equilibrium of the stressed calibration is an exact root.

3.4. From a point test to a validity interval, and completeness

The central step specialises right-hand-side ranging, the one-parameter case of multiparametric linear programming [15,16], to the objects of PvBT's Theorem 1, which tests a stored basis at a single demand level. The interval itself is standard. Its application to the linked problem is the contribution. Because $b(D)$ is affine, the basic solution $y_B(D) = u + D\,v$ (with u and v as defined in Section 3.2) is affine in D as well. Since dual feasibility does not depend on D, the sign condition $y_B(D) \ge 0$ is the only binding constraint on optimality, and the stored basis B is optimal exactly for

$$D \in \left[\, \max_{\ell:\, v_\ell > 0} (-u_\ell / v_\ell)\,,\; \min_{\ell:\, v_\ell < 0} (-u_\ell / v_\ell) \,\right], \tag{12}$$

where the bounds follow directly from u and v, with the conventions $\max \emptyset = -\infty$ and $\min \emptyset = +\infty$. When the LP tableau contains surplus variables (basic slacks that must satisfy $y_{B,\ell} \le 0$ rather than ≥ 0), each such row is included by substituting $(-u_\ell, -v_\ell)$ for (u_ℓ, v_ℓ), reducing the constraint to the same ≥ 0 form before applying (12).

Chaining these intervals from the lower end of a demand range enumerates every piece of the BU response g (subject to the degeneracy caveat of Appendix A), which is therefore an exactly known

[8] Nonsingularity of J_{CGE} at each solve ensures that f is locally single-valued along the Newton-tracked branch, making the scalar decomposition self-consistent. It does not rule out further interior CGE solutions off that branch (this is a named caveat in Appendix A), and Section 4 tests the branch by re-solving the CGE from a fresh benchmark start at a subsample of scan points.

piecewise rational function with exactly known breakpoints, obtained from the very objects the PvBT implementation already stores. Four consequences follow. First, the equilibrium condition (10) decomposes into finitely many smooth problems on known intervals, so that all roots in a range can be enumerated and classified by the sign of $|\varphi'| - 1$. Second, continuity of g, which the bracketing guarantee assumes, can be checked rather than assumed, by comparing the rational pieces at each breakpoint. Third, uniqueness, when it holds, is established up to the stated failure modes of Appendix A, in particular a root pair closer than the refinement grid inside one smooth cell. Roots are located from a fine grid whose cells (subintervals in ι, as distinct from the pieces of g in the demand D) are first split at the preimages of every piece breakpoint, found by bracketed root finding on the smooth CGE response, so that each scanned cell lies inside a single piece and h is smooth there.

Within each such cell, sign changes are polished, and interior extrema of h, detected by the analytic derivative (11), trigger a recursive split. The alignment matters because an unaligned grid would scan across the derivative jumps at the breakpoints, and the recursion matters because $h' = g'f' - 1$ need not be monotone within a piece even though the linear-fractional g keeps the sign of g' fixed on it. The fourth consequence is that the pieces are a property of the BU model alone. In an experiment design that varies only TD-side parameters, as the capital–labour shock grid of PvBT does, one enumeration serves every case, so the root finder evaluates g in closed form and the BU model is never solved again after the enumeration. The four steps of this section close the shortcomings of Section 2.3. Section 3.5 sets out how they relate to, and complement, the coordination algorithm of PvBT.

3.5. CBE: scope and complementarity with the coordination algorithm

Table 1 summarises the complementary roles of the coordination algorithm of PvBT and CBE. The two address distinct questions. PvBT reduce the cost of each bottom-up round, while CBE addresses the reliability and completeness of the answer. On convergence, the alternation requires an attracting fixed point, a condition that is untested in practice, whereas the bracketing solver of Section 3.1 requires only a sign change of h on a bracket and imposes no stability assumption. On stopping, the alternation halts when iterate change falls below a tolerance, certifying only that it stalled.

By contrast, the certification step applies a Newton–Kantorovich test to the integrated equilibrium system at the candidate, establishing that a true equilibrium exists nearby and is locally unique, or returning an explicit failure report. The test applies to any candidate, including a point returned by the alternation itself, so the certificate can be adopted without the bracketing solver. On scope, the alternation returns one equilibrium chosen by the starting point and gives no indication of others. The validity-interval enumeration of Section 3.4 returns all equilibria in a stated range and classifies each as attracting or repelling. On the use of the stored basis, PvBT use it for point prediction alone, whereas CBE also recovers the exact derivative (11) and the validity interval (12) from the same inverse, so the acceleration and CBE apply jointly rather than in competition.

4. Application to the CGE-TIMES linkage

The scheme of Section 3 is now applied to the CGE-TIMES instance of Section 2 as also employed by PvBT. All results are produced by the accompanying Python implementation (numpy only, with an internal revised simplex that returns optimal bases) and can be reproduced with the accompanying code. The subsections that follow are preceded by a validation of the implementation against the published GAMS results of PvBT, after which the equilibrium structure, computational cost and completeness of the results are examined in turn. The primary empirical result is the closed-form gain at which the corner equilibrium is born, and the margin it implies at the $s_K = s_L = 1.1$ point (Section 4.7). The stressed calibration of Section 4.2 is the vehicle for computing it, not a claim that multiplicity is probable at realistic parameterisations.

Table 1
Coordination algorithm versus certified bracketing and enumeration.

	Coordination algorithm [3]	CBE (this paper)
Question addressed	Cost of each round	Reliability and completeness of the answer
Convergence needs	Attracting equilibrium, $\|\varphi'(\iota^*)\| < 1$, untested	A sign change of h on a bracket; no stability condition
Stopping rule	Iterate change below a tolerance	Integrated-system violation $\lesssim 5 \times 10^{-11}$, and the candidate passes the Newton–Kantorovich test $2\beta L\eta \le 1$ with its ball inside the validity interval (12)
Output on termination	A point where the iteration stalled	A point at which a true equilibrium provably exists (up to floating-point evaluation, Section 3.3) and is locally unique, or an explicit failure report
Multiple equilibria	One, selected by the starting point; others undetected	All equilibria in a stated range, classified attracting or repelling, subject to the caveats of Appendix A
Repelling equilibria	Unreachable in principle	Computed routinely
Use of the stored basis	Point prediction of the BU solution (their Theorem 1)	Prediction, exact derivative (11), and validity interval (12)
Speed accounting	GAMS runtimes (their Table 1)	Solve counts in a single environment; no wall-clock claim versus GAMS

Note: The two approaches address complementary questions; CBE retains the per-round acceleration of the coordination algorithm. The parametric variant's cross-case saving assumes a fixed LP cost vector and constraint matrix across cases, a condition guaranteed within any single scenario run by the scalar linkage structure. Tolerances are worst cases over the 49 test instances (Section 4).

4.1. Replication and validation

Before any modification is evaluated, the implementation is validated against the published GAMS results. The base CGE reproduces the benchmark social accounting matrix (SAM) of [2] exactly, and every converged equilibrium of the capital–labour shock grid matches the GAMS reference, the small residual for the exact roots reflecting the heuristic stopping of the GAMS iterates rather than a discrepancy in the port. Iteration counts match for both the naive and the coordinated algorithm except where optimal-basis tie-breaks under degeneracy intervene.[9]

4.2. The round map and its slope

The exact piece structure of Section 4.6 makes the gain at which the corner becomes an equilibrium computable in closed form (Section 4.7). The stressed calibration examined below illustrates the equilibrium structure that lies beyond that threshold. Fig. 2 plots the round map φ for the published calibration ($\kappa = 1$, where κ scales the response of the bottom-up demand projections to the demand multiplier D) and for a stressed calibration ($\kappa = 20$) of the same model, in both cases at the $s_K = s_L = 1.1$ point of the shock grid. The gain multiplies the demand the bottom-up model receives, and the CGE never sees it, so the round map is $\varphi(\iota) = g\big(\kappa f(\iota)\big)$ and $h' = \kappa\, g'f' - 1$. Because g, its pieces, the derivative (11) and the validity interval (12) are all stated in the

[9] The reference is the GAMS output file BendersLink.lst, with the naive iteration counts additionally checked against Table 1 of [3]. On the 7×7 grid the initial LP objective is 52 237.0901 in both the port and the listing, the demand multiplier matches to four printed decimals and the BU objective to 2.3×10^{-7} relative (3.6×10^{-7} for the exact roots), and iteration counts match in 47 of the 49 cases.

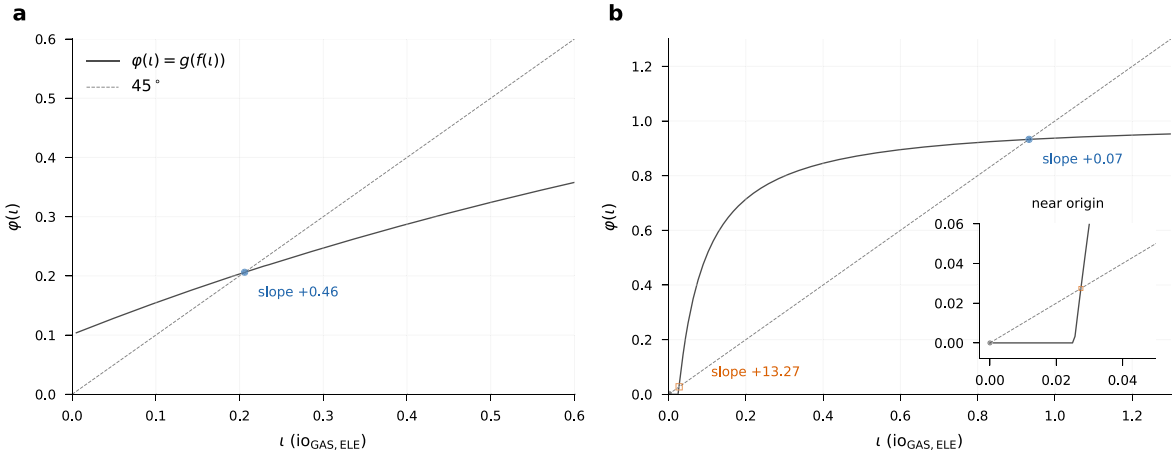

Fig. 2. Fixed-point structure of the round map at the $s_K = s_L = 1.1$ shock-grid point. The round map $\varphi(\iota) = g(\kappa f(\iota))$ against the diagonal for the published calibration ($\kappa = 1$, panel (a)) and the stressed calibration ($\kappa = 20$, panel (b), with a zoom near the origin). Equilibria are marked by stability: filled circles attracting, open squares repelling, each annotated with the map slope φ'.

demand the LP receives, the pieces do not depend on κ. A larger gain only carries a scan across more of them, and the displays of Sections 2 and 3 are the case $\kappa = 1$, where that demand is D itself. At $\kappa = 1$ the map crosses the diagonal once, with slope 0.46 at the crossing (the benchmark-point slope is 0.47, away from the crossing), so the local convergence condition $|\varphi'| < 1$ is satisfied and explains the favourable behaviour reported by PvBT. At $\kappa = 20$ the map instead crosses the diagonal at $\iota^* = 0.0274$ with slope 13.3 and at $\iota^* = 0.9329$ with slope 0.07, and the corner $\iota = 0$ is a third equilibrium, this one attracting, in which demand is low enough for hydro capacity to cover it. The two interior equilibria carry Newton–Kantorovich existence-and-uniqueness certificates, the test passing with the ball's demand excursion inside the basis validity interval at both. The corner is an exact root, with $h(0) = 0$ exactly, map slope 0, and residual zero. Interiority of the CGE is not delicate at this point. With $\iota = 0$ the gas sector loses its deliveries to electricity but keeps its other intermediate demands and household consumption, and the corner solve returns $X_{\mathrm{GAS}} = 6.64$ against a benchmark value of 10, with the smallest of the 22 unknowns, $L_{\mathrm{GAS}} = 0.66$, well away from zero. Its Newton–Kantorovich ball reaches past an endpoint of the validity interval of the basis covering it, so the requirement of Section 3.3 that the ball's demand excursion lie within that interval is not met and no certificate is issued. The Kantorovich inequality itself is satisfied, and the corner is an exact root, so what is lacking is a certified uniqueness radius rather than existence. The middle equilibrium is repelling and separates the gas-free power regime from the gas-intensive one. The enumeration classifies these equilibria by stability only. Which equilibrium an analyst should report is a selection question the certified set makes explicit without answering, and a welfare ranking across the three, computable from the stored CGE solutions, is not pursued here.

4.3. Path dependence of the alternation

Panel (b) of Fig. 3 traces the alternation at $\kappa = 20$ from three starting points. From $\iota_0 = 0.02$, below the repelling equilibrium, the iteration collapses to the corner in two rounds. From $\iota_0 = 0.0286$ or from the benchmark value 0.1818 it converges to the gas-intensive equilibrium. No tested starting point yields the repelling equilibrium, no damping factor can stabilise it (Section 2.3), and no run produces any indication that other equilibria exist. The safeguarded scalar solver locates both interior equilibria, with fixed-point residuals below 2×10^{-11} and integrated-system violations below 5×10^{-11} at the repelling and below 2×10^{-10} at the gas-intensive equilibrium, and each carries a Newton–Kantorovich existence-and-uniqueness certificate.

This equilibrium structure, in which two equilibria are reachable by alternation and one is not, is the setting against which the computational cost of each method is compared.

4.4. Work comparison

Panel (a) of Fig. 3 compares the number of full-model solves required for the 49-case grid. Every total is counted directly at the solver call sites, with the calibration run and the per-case initialisation solves included once each, exactly as they occur. The naive alternation uses 818 CGE and 818 LP solves. The PvBT coordination algorithm uses the same 818 CGE solves and reduces the LP count to 133 on the accounting used here, reproducing the benefit they report. The safeguarded scalar solver needs 367 CGE and 144 LP solves with Brent's method, and 347 and 167 with the Newton variant, in both cases including the bracketing phase and the certification solves, and without any warm state carried between cases.

The scalar solver therefore buys certified output and independence from contraction with fewer than half the TD solves, at a BU effort slightly above that of the coordination algorithm. One caution about that comparison concerns the motivating setting of PvBT, in which the BU model is the expensive one, so measured against their own cost model the scalar solver alone is no improvement (144 and 167 BU solves against the coordination algorithm's 133), and the halved TD count matters only where the TD model is costly. In a full-scale TIMES deployment, where the LP is by far the more expensive of the two models, those additional 11 to 34 BU calls represent a material cost that the halved TD count does not offset.

However, this comparison of solve counts is secondary to the question CBE is designed to answer. The certified equilibria, stability classifications, and multiplicity counts of Sections 4.2–4.7 are products of the within-run enumeration and hold for every scenario regardless of whether A and c are shared across cases. The additional BU-cost reduction of the parametric variant is conditional on that scenario structure.

The case on PvBT's own terms is made by the parametric root finder of Section 3.4, which removes the bottom-up cost almost entirely. Since the shocks only enter the CGE, one enumeration of the validity

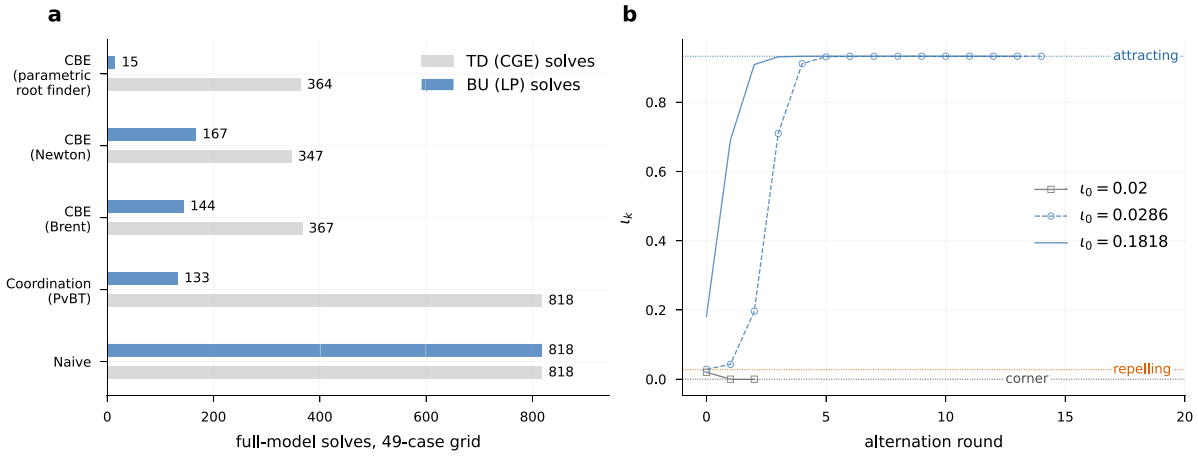

Fig. 3. Computational work and path dependence of the alternation. Panel (a) Full-model solves needed for the 49-case shock grid by the naive alternation, the coordination algorithm of [3], the CBE scalar solver in its Brent and Newton variants, and the CBE parametric root finder, which shares one piece enumeration across cases (certification included throughout, all totals counted at the solver call sites). Panel (b) Alternation trajectories at $\kappa = 20$ from three starting points, against the three equilibria (dotted lines). The repelling equilibrium is unreachable by alternation from any tested start.

intervals serves all 49 cases,[10] and the grid solves with about 364 TD (CGE) and 15 BU (LP) solves in total, certificates included (assembled from the stored bases by arithmetic, not by further solves). Worst-case certificates are consistent across all certified variants, with integrated-system violations below 5×10^{-11} in every case on the 49-instance grid, and the Newton–Kantorovich existence test passing at every case with worst-case $2\beta L\eta$ below 10^{-6}, existence radius below 2×10^{-10}, and uniqueness ball of radius at least 4×10^{-4}.

The scope of this comparison requires clarification. Wall-clock times of the present implementation are not comparable to the runtimes reported in Table 1 of [3], which reflect a different language, solver stack and machine, and no such comparison is made here. The comparison is between algorithms, all re-implemented and run in the same environment, with work measured in counts of full model solves. Counts give an implementation-independent proxy, but not a perfect one. They are also not bitwise reproducible, since under the same LP degeneracy that produces the iteration-count tie-breaks of Section 4.1, a tie in the optimal basis can change a few pivots, so a solve count shifts by a few between runs (and the enumeration sweep of Section 4.6, which scans tens of thousands of times, by a few hundred). Every solve total quoted in this paper is therefore representative of a typical run rather than an exact invariant, but the certified results accompanying each quoted total (the equilibria, residuals and uniqueness counts) are reproducible. A CGE solve and an LP solve have different costs, and within a given solver a warm-started re-solve is cheaper than a cold solve, so the totals in Fig. 3 indicate the structure of the saving (half the TD solves for the scalar solver, and almost all of the BU solves for the parametric variant, certification included in both) rather than a precise speed ratio. The qualitative differences, certified output and independence from contraction, do not depend on this accounting.

[10] The condition that makes one enumeration sufficient, namely that A and c are fixed across all 49 cases and only b varies through the affine channel $b(D) = b_0 + D\,\Delta b$, is the cross-case form of the condition on which PvBT's sign test rests within a case. The test is sufficient for optimality because a stored basis carries a dual certificate, the non-negativity of reduced costs, that is preserved whenever c is fixed (Section 2). When technology costs or LP structure vary across scenarios, the intervals require re-enumeration for CBE, so the parametric variant operates per scenario. CBE's within-scenario contributions, certified equilibria, multiplicity detection, and stability classification, follow from the within-run enumeration and hold for every scenario regardless of how scenarios differ from one another.

4.5. The integrated solve as a baseline

Several routes to the linked equilibrium can be distinguished. The integrated formulation of Böhringer and Rutherford [7], which rests on Mathiesen's complementarity treatment of economic equilibria [4], solves the top-down and bottom-up models as one mixed-complementarity problem.[11] The remaining routes decompose, but in two distinct senses. Böhringer and Rutherford themselves later decompose that integrated system [12], splitting it into an economic and an energy block whose iteration converges back to the same integrated equilibrium, whereas the coordination algorithm of PvBT [3] and the CBE scheme of this paper proceed from the soft-linking alternation between two independently formulated models, solving each separately and passing information between them. That decomposition, which Lanz and Rausch [27] apply at scale to electricity-generation technologies, is itself a fixed-point iteration between blocks, so the convergence and stopping questions of Section 2.3 apply to it as they do to the soft-linking alternation, and its guarantees are of the same conditional kind. PvBT stay within the alternation they inherit and lower the cost of each bottom-up round by predicting LP solutions from stored optimal bases, drawing on that integrated line [7] only to establish that the top-down block is small relative to the round. CBE reduces the linkage to a scalar root, and so decomposes as well, yet it puts the integrated formulation to two further uses. The residual of the full integrated system certifies every candidate equilibrium independently of the path that produced it, and a direct damped-Newton solve of the same system is the baseline against which the iterative schemes are measured below.

That baseline does not iterate at all. The implementation runs a damped Newton method on the CGE unknowns and ι jointly, with the LP represented by its current optimal basis (so the linking equation is the rational function of Section 3.4), obtained by one LP solve per case and re-solved only when the sign test of PvBT's Theorem 1 rejects that basis. Applied to the 49-case grid, the totals are 111 LP solves and 208 Newton iterations on the stacked system, each iteration costing one factorisation of the 23×23 Jacobian of that stacked system, a fraction of a full CGE solve, plus the same per-case certification as the other methods (49 CGE and 49 LP solves).

[11] This comparison assumes that the integrated system can be assembled and solved as one object, which the small test model permits but full deployments often do not. Where the bottom-up model is large, separately maintained, or available only as a solver rather than in closed algebraic form, the single-system route is closed and the coupling must proceed by decomposition.

All 49 results match the GAMS reference, with integrated-system violation below 5×10^{-11}. This comparison cuts both ways, as Fig. 3 reports.[12] Measured by TD-side work, integration dominates every iterative scheme here, with 208 cheap iterations against hundreds of full CGE solves. Measured by PvBT's own BU-heavy cost model it does not dominate, since its 160 LP solves with certification (111 without) exceed the 133 of the coordination algorithm, which certifies no equilibrium, and the parametric variant's 15 remains the best BU count by an order of magnitude. What integration cannot supply, regardless of computational effort, is the global picture. At $\kappa = 20$ the integrated Newton converges to the gas-free corner from every tested start, including $\iota_0 = 0.0274$ immediately adjacent to the repelling equilibrium, and, like the alternation, gives no indication that other equilibria exist. The contribution of Section 3.4 is accordingly not the cost of one solve but completeness and the stability margin of Section 4.7, which no single-point method provides. The exact piece structure on which that completeness rests is examined in the following subsection.

4.6. Exact pieces and completeness

Fig. 4 shows the BU response g on $D \in [-0.15, 0.6]$, computed from the validity intervals of (12). Fifteen bases cover this range, and the closed-form pieces agree with direct LP solves at random demand levels to 5×10^{-16} in the activity ratio (the LP objective to 3×10^{-11} in absolute terms). With the pieces in hand, complete enumeration over $\iota \in [0, 1.5]$ gives the following. At $\kappa = 1$, every one of the 49 published cases possesses, on the tracked top-down branch, exactly one equilibrium, which coincides with the GAMS solution.

Because the shocks only enter the CGE, this enumeration sweep shares a single piece store across cases and costs about 37 600 CGE solves and 63 LP solves in total over 29 stored pieces.[13] The sweep is a one-time offline calculation whose piece store is reused across all 49 scenario cases and underpins the closed-form margin of Section 4.7. The sweep establishes uniqueness across the whole range rather than locating a single equilibrium, and is correspondingly more costly than the parametric root finder of the work comparison. Most of the CGE total goes on evaluating h at the grid points and its analytic derivative at the ends of every scanned cell. The scan points inserted at the piece-breakpoint preimages so that every scanned cell lies within a single piece (Section 3.4) are a minor share. At a subsample of scan points in each case the warm-started CGE branch is also re-solved from a fresh benchmark start, with the largest deviation being 2.5×10^{-13}, so the branch tracking that underlies the scalar reduction introduces no hysteresis on this model. Subject to the grid-resolution and branch

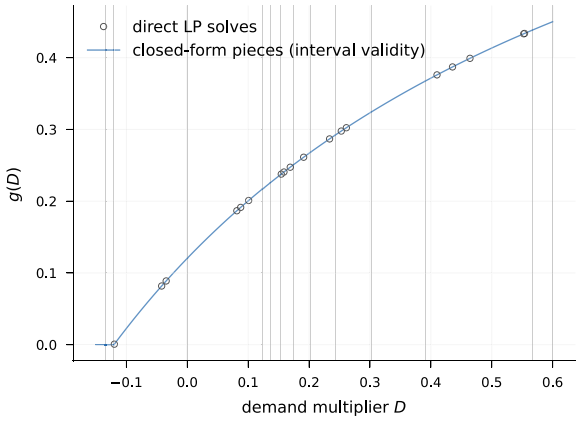

Fig. 4. Exact piecewise structure of the bottom-up response. $g(D)$ computed from basis validity intervals (solid line, breakpoints as vertical lines), verified against direct LP solves at random demand levels (open circles). Agreement is at machine precision.

caveats of Appendix A, the published results are therefore not only converged but unique in the stated range. At $\kappa = 20$ the one scan returns all three equilibria of Fig. 2, the corner included.

This corner is an exact root on the scan grid. The implementation re-scans the adjacent cell from just inside, shrinking the offset under the analytic slope until the near side of any neighbouring root is located, so equilibria near the corner are recovered down to a separation of 10^{-7} in ι, below which the enumeration treats two roots as one. The scan finds no jump at any of the 45 interior breakpoints of the 46 pieces of this $\kappa = 20$ scan, verifying the continuity hypothesis of the bracketing guarantee for this model rather than assuming it, and classifies each equilibrium by its analytic slope. Residuals at all enumerated equilibria are below 2×10^{-14} and integrated-system violations below 2×10^{-10}, and every interior enumerated root carries a Newton–Kantorovich existence-and-uniqueness certificate.

The availability of exact pieces makes it possible to compute, in closed form, the gain at which the corner equilibrium is born.

4.7. Distance to corner birth

The stressed calibration raises the question of how far the published model is from equilibrium multiplicity. With the exact pieces available, the answer is closed-form rather than purely numerical. Recall that κ is the gain on the demand link of Fig. 2, equal to one at the published calibration. The CGE never sees κ, so the demand level attached to the corner $\iota = 0$ is $f(0)$, one CGE solve, and the LP receives the effective demand $\kappa f(0)$. On the piece store the gas-power activity is exactly affine in demand and vanishes at the breakpoint D_{brk} that separates the pieces on which gas-fired power is idle from those on which it runs. The corner therefore becomes an equilibrium exactly at the gain $\kappa^* = D_{\text{brk}}/f(0) = 5.5690$, with $D_{\text{brk}} = -0.1205159$ and $f(0) = -0.0216403$ at the $s_K = s_L = 1.1$ point. At this point the repelling threshold detaches into the interior, a boundary crossing rather than an interior fold. Whether corner birth coincides with the onset of multiplicity, or instead absorbs the incumbent equilibrium, which then sits at the corner and remains attracting, is determined numerically case by case. Two checks accompany the value. Prediction-free evaluations confirm $h(0) > 0$ just below κ^* and $h(0) = 0$ just above, and an independent bisection on the equilibrium count over $\iota \in [0, 1.5]$, with the corner counted, brackets $\kappa^* \in (5.5625, 5.5703]$, which contains the closed-form value.

The two sides of κ^* carry different kinds of certainty. Above it the corner equilibrium exists for every gain, not only the probed ones, because the effective corner demand $\kappa f(0)$ moves monotonically

[12] The integrated Newton baseline of Section 4.5 is not charted because its 208 stacked-Newton iterations cost approximately one CGE Newton iteration each rather than one full CGE solve, so a bar on the same axis would misrepresent the effort. Its full count is 111 LP solves plus 208 iterations plus 49 CGE and 49 LP certification solves.

[13] The CGE solve total is the budget of the one-time enumeration that underpins Section 4.7, not a per-scenario cost. Two observations bear on whether this is prohibitive at realistic model scales. First, the scan admits full parallelisation across its scan points, since each is an independent evaluation of $h(\iota)$ requiring one CGE and at most one LP solve with no dependency between points. Wall time therefore scales with the number of available processors rather than with the total solve count, and on a high-performance cluster the enumeration is likely to remain feasible even for large-scale TIMES deployments. Second, the number of pieces is a property of the LP's constraint matrix and cost vector, and the parametric path $b(D) = b_0 + D\,\Delta b$ moves along a single direction in right-hand-side space, driven only by the demand rows of the BU model. The piece count therefore scales with the demand-row structure rather than with total model size. Heavy LP degeneracy could inflate this through cascading pivots at breakpoints, so the empirical scaling of the enumeration to full-scale deployments remains an open question this paper does not address.

through the breakpoint, and a repelling threshold separating it from any coexisting interior equilibrium follows by continuity. Below it, uniqueness is certified at the probed gains by complete enumeration, each subject to the caveats of Appendix A. Starting the count scan at $\iota = 0$ also matters, since a scan floored at a small positive value records instead the gain at which the newly born threshold clears the floor, and overstates κ^* by about 0.01 already at a floor of 10^{-4}. The $\kappa = 20$ of Fig. 2 is therefore not the onset of the phenomenon but a point chosen beyond it for visual clarity, and at the $s_K = s_L = 1.1$ point the model retains a unique equilibrium until the demand-link gain reaches about 5.6, a margin of about 4.6 above its calibrated value of one. The margin is case-specific. Since $f(0)$ varies across the shock grid, κ^* is defined at 19 of the 49 cases and ranges from 1.15 to 25.9, while at the remaining 30 the corner route cannot fire at any gain. At 16 of the 19 the equilibrium count passes from one to three at κ^*, so corner birth is the onset of multiplicity there. The corner instead absorbs the incumbent at the three smallest values, which belong to the $s_K = s_L = 1$ point and to $(s_K, s_L) = (1.05, 1)$ and $(1, 1.05)$. Uniqueness at those three points then holds until an interior fold at about 3.5, 2.8 and 2.5. Just below the smallest onset value, 1.69 at the $s_K = 1.1$, $s_L = 1.0$ point, every published case still has exactly one equilibrium, as it does at gain one.

Two further measurements locate where that margin resides. The contraction margin $|\varphi'|$ at the 49 published equilibria lies in $[0.331, 0.592]$, and it is nearly invariant to the CES substitution elasticities. Recalibrating the CGE with σ_i scaled by factors of 1.5 to 3 (the SAM pins the benchmark value shares, so the elasticities are the standard free parameters) moves the maximal slope over a 25-point sample of $\iota \in (0, 1]$ at the $s_K = s_L = 1.1$ point from 0.557 only to 0.558. In this instance, the feedback strength of the linkage is set by the demand-projection formula and the slopes of the LP pieces, not by factor substitution.

5. Discussion

This section asks what the guarantees of Sections 3 and 4 rest on and where they stop. It names the four edges of the envelope within which they hold and closes on the question of scale that the envelope leaves open. Section 6 then concludes.

The single structural change, from seeking a fixed point to seeking a root, carries the convergence guarantee, the certificate and the enumeration. A fixed-point iteration converges only under contraction, stops on a heuristic, and reaches one equilibrium chosen by its starting point, whereas a root admits a path-independent certificate and, once the response is known in closed form on each interval, every equilibrium in the range. These achievements depend only on A, c, and the affine right-hand-side structure holding within a single run, a condition the scalar linkage guarantees by construction. None of it displaces the per-round acceleration of PvBT, which CBE retains and reuses, since the stored basis inverse that predicts a skipped solve also supplies the exact derivative and the validity interval.

Because these guarantees hold within a stated envelope, naming its edges is part of stating the result. The first edge is the scalar exchange. This restriction is shared with the PvBT setting. The coordination algorithm is stated for the general form of (1)–(2) but demonstrated, and its computational savings measured, on the same scalar-linked instance. The reduction to a scalar root and the completeness argument both rely on the linkage passing one scalar in each direction, as in the PvBT experiments and in linkages of the Helgesen type (one technology coefficient against one demand multiplier). With vector-valued exchange the validity intervals of Section 3.4 become polyhedral critical regions, still exactly computable in the sense of multiparametric programming [15] but rapidly more numerous, and bracketing has no direct analogue in higher dimension, where guaranteed fixed-point computation belongs to the simplicial and homotopy methods in the tradition of Scarf [28]. The certification step, by contrast, carries over to any linkage dimension, since evaluating the integrated equilibrium residual at a candidate is always cheap. The Newton–Kantorovich bound on $\|J^{-1}\|$ and the Lipschitz constant, however, grow in cost with dimension.

The second edge is the linear programme. PvBT operate under the same restriction. Their Theorem 1 relies on the LP optimal basis in precisely the way that CBE does. The stored optimal basis supplies both the exact derivative and the validity interval, and right-hand-side ranging makes the BU response an exactly known piecewise rational function. When the BU model is a mixed-integer programme (MIP), it has no optimal basis of this kind. Within a fixed integer assignment the block reduces to an LP and the response is still piecewise rational, but the optimal value function in the demand multiplier is piecewise linear, non-convex and generally discontinuous, so the response jumps where the optimal integer assignment switches and the bracketing guarantee, the completeness argument and the cheap certificate all lapse across those switches.[14] Energy-system models used in practice frequently carry integer features such as lumpy investment or unit commitment, so this is a bound on where the guarantees hold rather than a defect within the stated scope.

The third edge is the direction of the linkage. Throughout, the bottom-up cost vector is fixed, so the energy system responds to the economy through quantities while its marginal costs never reprice the CGE (Section 2). Soft links in the TIMES-CGE literature often run a price channel as well, returning bottom-up marginal costs to the economy or feeding economy-side fuel and carbon prices into the bottom-up objective [2,9,10]. The first direction exchanges a dual quantity that the present scalar reduction does not treat, and the second makes c depend on the top-down solution, which invalidates the stored dual certificates and with them the validity intervals, the exact derivative and the parametric enumeration. The guarantees of this paper are quantity-linked by construction, and extending them to price-linked exchange is an open direction.

The fourth edge is scale. The PvBT model employed in this paper is small, and the multiplicity exhibited under the stressed calibration arises under an amplified demand link rather than at the published calibration, whose margin at the illustrated point follows in closed form from the piece structure. The finer numerical caveats that bound the enumeration within this envelope, on grid resolution, degeneracy and cross-solver reproducibility, are collected in Appendix A.

Whether full-scale linkages cross the multiplicity margin at plausible calibrations is an empirical question this paper does not settle. In applied general equilibrium, uniqueness has only been established model by model [22], and the validity-interval machinery is what makes the same question answerable for the linked energy-economy system. Within the scalar-exchange envelope the enumeration answers it completely, and beyond that envelope the certification step still carries over to any linkage dimension while the contraction margin remains meaningful wherever a round map exists.

[14] The lapse is specific to each component. MIP solvers do not return sensitivity ranges or dual variables after a branch-and-bound solve, so the validity interval (12) and the exact derivative (11) are unavailable regardless of how the solver is accessed. The LP-based components of the certificate (reduced-cost non-negativity and complementary slackness) are likewise absent, leaving only a feasibility check on the primal solution against the linking equations and the CGE conditions. A fix-and-reoptimise step, which holds the integer variables at their optimal values and re-solves the resulting LP, recovers a local basis and dual certificate at the cost of one additional solve per iterate, restoring local derivatives but not the global validity interval or completeness. Continuity of g, which the bisection guarantee requires, is not assured when the optimal integer assignment changes with D, since the switching points are not computable in closed form. The practical alternatives are bisection on a fine grid with explicit discontinuity detection, surrogate modelling of g from a batch of MIP solves, or the fix-and-reoptimise strategy applied at each candidate iterate.

6. Conclusion

This paper adds reliability and completeness to the accelerated top-down and bottom-up coordination of PvBT. The scalar-linked equilibrium becomes the root of a scalar function, and certified bracketing locates that root without any contraction assumption. The pointwise basis test of PvBT is extended to closed-form validity intervals, under which the bottom-up response is an exactly known piecewise rational function and every equilibrium in a stated range, on the tracked top-down branch, can be enumerated and classified. On the CGE-TIMES instance of PvBT the scheme reproduces the published results and certifies that a true equilibrium exists at each and is locally unique. The published calibration has a single equilibrium in the relevant range, and a stressed calibration has three.

One recommendation carries beyond the model employed in this paper. The contraction margin $|\varphi'|$ is available at every certified solution at no additional cost, and reporting it alongside the solution is a cheap discipline that records whether a converged point is firmly attracting or close to the boundary at which the alternation would fail. That boundary is not hypothetical. Under an amplified demand link the same model acquires a repelling equilibrium that no damped alternation can reach, and the gain at which it is born has a closed form built from the validity intervals, so each calibration on which the corner route is open carries a computable margin to that birth, which the enumeration then classifies as the onset of multiplicity or as absorption of the incumbent.

The stated limits on scalar exchange, the linear programme, the quantity-only linkage, and scale each open a direction for further work. Beyond these, the integrated residual and the stored basis inverse are the objects that a certified Benders decomposition for the general TD-BU linkage would require, and developing that connection is a direction this paper does not pursue.

CRediT authorship contribution statement

Sergio Bobka: Writing – review & editing, Formal analysis. **Julia Gershenzon:** Writing – review & editing, Formal analysis. **Stefan Vögele:** Funding acquisition. **Jochen Linßen:** Funding acquisition. **Andrew G. Ross:** Writing – review & editing, Writing – original draft, Methodology, Investigation, Funding acquisition, Formal analysis, Conceptualization.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgements

This work was supported by the Helmholtz Association under the programme “Energy System Design”, and the RESUR project in phase II (grant number 03EI1123A) supported within the scope of energy research by the Federal Ministry for Economic Affairs and Energy on the basis of a decision by the German Bundestag. Open Access is funded by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) (491111487). The authors used Anthropic’s Claude as a language and code assistant during the preparation of this work. After using this tool, the authors reviewed and edited the content as needed and take full responsibility for the content of the publication.

Appendix A. Numerical caveats and reproducibility

The completeness and replication results of Sections 3.4 and 4 carry a small number of qualifications, set out here with their remedies. Three of them bound the completeness claim, concerning in turn the treatment of the top-down model, the resolution of the root enumeration, and the handling of degenerate pieces. None of the three is triggered by any run reported in the paper, yet each is a point at which a stronger guarantee would need additional machinery. A fourth qualification concerns the reproducibility of the iteration counts across solvers, and a closing note concerns scenario scope.

The first condition concerns the top-down model. The CGE is solved here as an interior square system whose interiority is verified at every solve, and a complementarity solver would be required only if boundary equilibria were of interest. The completeness statement also treats the CGE response f as a function, and the branch on which it does so can be stated once. It is the implicit-function-theorem continuation, in ι over the scanned range $[0, 1.5]$, of each case’s CGE solution from its benchmark start, extended through each solve by the nonsingularity of J_{CGE} (Section 3.2) and followed numerically by the warm-started Newton method, and the enumeration lists the equilibria on this branch. Multiplicity of interior CGE solutions at a fixed ι, of the kind general equilibrium theory cannot exclude [21], would add branches that the scan does not visit, although uniqueness is typical for calibrated models of this class [22]. The cold-start cross-check of Section 4 tests this branch assumption at sampled points, with a largest deviation of 2.5×10^{-13} on this model, but does not prove it globally. Within each certified equilibrium’s Newton–Kantorovich uniqueness ball a second branch is provably excluded, so the failure mode is discharged locally around every certified equilibrium. Branches outside every ball remain unvisited, and the global statement is unchanged.

The second concerns the resolution of the root enumeration. Roots are sought on cells aligned with the piece breakpoints and refined by analytic derivatives. The scan places 160 base points on $\iota \in [0, 1.5]$ (the corner $\iota = 0$ is itself a grid point, handled as described in Section 4.6), cuts every cell at the preimage of each piece breakpoint to a tolerance of 10^{-11}, and locates each root to a tolerance of 10^{-12}. The cuts presuppose a CGE response f that is monotone across the grid, which the validation suite asserts for every case reported here. The one configuration the scan could miss is then a pair of roots closer than this refinement grid inside a single smooth cell, which only interval arithmetic on the CGE side would exclude rigorously. Two devices narrow this mode without closing it. A pair of roots inside a smooth cell forces an interior extremum of h between them, so the derivative-based splitting of Section 3.4 detects the pair when h' changes sign across the containing cell and the extremum is located within four levels of recursive splitting, although a cell on which h' carries the same sign at both endpoints can still conceal one. The Newton–Kantorovich ball excludes a further root within r_+ of each certified candidate in the stacked unknowns, which constrains where a missed pair could sit but says nothing about the remainder of the range.

The third concerns degenerate pieces. The enumeration crosses degenerate breakpoints by a fixed step of 10^{-8} in the demand multiplier, probes every crossed gap by one direct LP solve, and stops with an error if the closed-form value deviates from that solve by more than 10^{-9}. The step sits well above double-precision rounding (about 2×10^{-16}) and well below the typical width of a piece, of order 10^{-2} across the demand range, so only a degenerate coincidence narrower than 10^{-8} can be skipped. The enumeration guarantee is therefore conditional on no piece being narrower than the 10^{-8} step while astride neither the probe point nor the gap edges. Lexicographic pivoting would remove the condition entirely, and no run reported in this paper crosses such a gap.[15]

The fourth concerns cross-solver reproducibility. Two of the 49 replicated iteration counts differ from the GAMS reference through

[15] This is distinct from the paired-roots condition above. There the risk is two roots inside an enumerated piece. Here it is a piece too narrow to enter the enumeration. The two are remedied differently, by interval arithmetic on the CGE side and by lexicographic pivoting respectively.

optimal-basis tie-breaking under degeneracy, which indicates the degree of agreement that any basis-dependent method, including that of PvBT, can attain across solvers.

Finally, one note concerns scenario structure. When scenario variation affects the LP cost vector or constraint matrix rather than the right-hand side alone, cross-scenario basis reuse is unavailable to the parametric variant of CBE, which then operates per scenario. CBE's within-scenario contributions are unaffected.

Appendix B. Nomenclature

For the notation used throughout the paper, this appendix provides a single reference. Symbols are grouped by block, with the top-down symbols following the source model of [2] as used by PvBT and the certification symbols following the standard Newton–Kantorovich statement [24]. Subscripts carry six distinct roles, namely an index over S (P_i), an LP row or basis position (u_ℓ, v_ℓ), a fixed label (P_K, $r_\pm$, b_0, $\|\cdot\|_\infty$), a name (J_{CGE}, D_{brk}), a process-period pair ($\mathrm{ACT}_{p,t}$), and an iteration or cut counter (ι_k, λ_k). Superscripts and accents carry roles of their own. An asterisk marks a value at a linked equilibrium (ι^*, x^*, y^*), and on κ alone the critical gain. A raised zero marks the benchmark (X_i^0), and a prime a derivative (f', φ'). A hat marks a prediction ($\hat{y}_B$) and a tilde a candidate ($\tilde{\iota}$, $\tilde{w}$), while a bar marks an endowment ($\bar{K}$, $\bar{L}$), a fixed value of x at which a basis is computed ($\bar{x}$), or, paired with an underbar, a bracket endpoint ($\underline{\iota}$, $\bar{\iota}$).[16]

Top-down (CGE) block	
S	Sector set {GAS, ELE, MAN, NON}
i, j	Sector index, $i, j \in S$
P_K	Return to capital
P_L	Wage, fixed as numeraire
P_i	Commodity prices
X_i	Gross output
X_i^0	Benchmark (calibration) gross output
$\bar{K}, \bar{L}$	Capital and labour endowments
K_i, L_i	Factor demands, L_i not the Lipschitz constant L
C_i	Consumption, distinct from the computed inverse C
R	Income
α_i	Marginal expenditure share, LES
μ_i	Subsistence consumption, LES
γ_i	Capital weight, CES value added
σ_i	Substitution elasticity, CES value added
a_i	CES scale parameter
Φ_i	CES factor-price aggregate raised to $1-\sigma_i$
io_{ij}	Leontief technical coefficients; ι is one entry
s_K, s_L	Multiplicative factors scaling $\bar{K}, \bar{L}$ across the shock grid
x	Top-down variable vector
$\bar{x}$	A fixed x at which a basis is computed
z	The 22 CGE unknowns, P_L excluded as numeraire
F	Top-down complementarity map, (1)
Bottom-up (LP) block	
y	Bottom-up variable vector, (2)
y_B, y_N	Basic and non-basic parts of y
$\hat{y}_B$	Predicted basic solution $B^{-1}b(x)$
c	LP cost vector
A	LP constraint matrix
n	Number of bottom-up variables
$b(\cdot)$	LP right-hand side, $b(x) = b(D(x))$ by (7)
b_0	Right-hand side at $D = 0$
Δb	Right-hand-side slope, $\mathrm{d}b/\mathrm{d}D$
B	Optimal basis
u, v	$B^{-1}b_0$ and $B^{-1}\Delta b$
ℓ	LP row or basis position index, (12)
λ_k	Optimal dual vector of stored basis k
θ, ζ_k	Cut-selection scalar and weights in the top-down master
$\mathrm{ACT}_{p,t}$	Activity level of process p in period t
Linkage	
D	Demand multiplier passed by the CGE, (7)
ι	$\mathrm{io}_{\mathrm{GAS,ELE}}$, the coefficient passed back, (8)
f	Half-round map $\iota \mapsto D$
g	Half-round map from the demand the LP receives to ι
φ	Round map $\iota \mapsto g(\kappa f(\iota))$
φ'	Round-map slope, with $\lvert\varphi'\rvert$ the contraction margin
h	$\varphi(\iota) - \iota$, whose roots are the linked equilibria, (10)
ι^*	Value of ι at a linked equilibrium, as x^*, y^* for x, y
$\underline{\iota}, \bar{\iota}$	Bracket endpoints for root finding
ω	Damping factor of the alternation
κ	Gain applied to D before the LP, unity at calibration
κ^*	Gain at which the corner becomes an equilibrium
D_{brk}	Piece breakpoint at which gas-fired generation vanishes
Certification	
Ψ	Integrated equilibrium residual, distinct from F
w	The 23 unknowns (z, ι)
$\tilde{\iota}, \tilde{w}$	Candidate equilibrium, in ι and in stacked form
J	$\partial\Psi/\partial w$, the 23×23 stacked Jacobian
J_{CGE}	The 22×22 CGE Jacobian
s	Derivative of the CGE residual in ι, distinct from s_K, s_L
C	Computed inverse of J, as in $\|I - CJ\|_\infty < 1$
β	Verified upper bound on $\|J^{-1}\|_\infty$
η	Verified upper bound on the first Newton step length
L	Certified enclosure of the Lipschitz constant of J
r_-, r_+	Existence and uniqueness radii
General notation	
$\perp$	Complementarity, (1)
$(\cdot)^\mathsf{T}$	Transpose
I	Identity matrix
$\|\cdot\|_\infty$	Maximum norm, on vectors and on matrices
cond_∞	Condition number in $\|\cdot\|_\infty$

[16] The CGE is calibrated to the benchmark social accounting matrix of [2] and its quantities are therefore in the value units of that matrix, with P_L fixed as numeraire and the remaining prices normalised against it. The bottom-up activities $\mathrm{ACT}_{p,t}$ are in the energy units of the TIMES instance. The two exchanged scalars D and ι are dimensionless ratios, as are κ, s_K, s_L and φ'.

Data availability

Data and model code are available at https://github.com/ag-ross/TDBU. Some small code snippets are adapted from [29].

References

[1] Loulou R, Labriet M. ETSAP-TIAM: the TIMES integrated assessment model Part I: Model structure. Comput Manag Sci 2008;5:7–40. http://dx.doi.org/10.1007/s10287-007-0046-z.

[2] Helgesen PI, Tomasgard A. From linking to integration of energy system models and computational general equilibrium models – Effects on equilibria and convergence. Energy 2018;159:1218–33. http://dx.doi.org/10.1016/j.energy.2018.06.146.
[3] Pisciella P, van Beesten ER, Tomasgard A. Efficient coordination of top-down and bottom-up models for energy system design: An algorithmic approach. Energy 2023;284:129320. http://dx.doi.org/10.1016/j.energy.2023.129320.
[4] Mathiesen L. Computation of economic equilibria by a sequence of linear complementarity problems. In: Manne AS, editor. Economic equilibrium: model formulation and solution. Berlin, Heidelberg: Springer Berlin Heidelberg; 1985, p. 144–62. http://dx.doi.org/10.1007/BFb0121030.
[5] Dixon P, Rimmer M. Small-group monopolistic competition in a global computable general equilibrium model: Meeting the Markusen challenge. J Glob Econ Anal 2025;10(2):46–65. http://dx.doi.org/10.21642/JGEA.100202AF.
[6] Hourcade J-C, Jaccard M, Bataille C, Ghersi F. Hybrid modeling: New answers to old challenges. Introduction to the special issue of the energy journal. Energy J 2006;27:1–11. http://dx.doi.org/10.5547/ISSN0195-6574-EJ-VolSI2006-NoSI2-1.
[7] Böhringer C, Rutherford TF. Combining bottom-up and top-down. Energy Econ 2008;30(2):574–96. http://dx.doi.org/10.1016/j.eneco.2007.03.004.
[8] Wene C-O. Energy-economy analysis: Linking the macroeconomic and systems engineering approaches. Energy 1996;21(9):809–24. http://dx.doi.org/10.1016/0360-5442(96)00017-5.
[9] Krook-Riekkola A, Berg C, Ahlgren EO, Söderholm P. Challenges in top-down and bottom-up soft-linking: Lessons from linking a Swedish energy system model with a CGE model. Energy 2017;141:803–17. http://dx.doi.org/10.1016/j.energy.2017.09.107.
[10] Fortes P, Pereira R, Pereira A, Seixas J. Integrated technological-economic modeling platform for energy and climate policy analysis. Energy 2014;73:716–30. http://dx.doi.org/10.1016/j.energy.2014.06.075.
[11] Elberry AM, Fragkiadakis K, Paroussos L, J. van Stralen M Scheepers, Sijm J, Faaij A, van der Zwaan B. Soft-linking a general equilibrium model and an energy system model: towards a carbon–neutral economy by 2050. Energy Convers Manag: X 2026;30:101704. http://dx.doi.org/10.1016/j.ecmx.2026.101704.
[12] Böhringer C, Rutherford TF. Integrated assessment of energy policies: Decomposing top-down and bottom-up. J Econom Dynam Control 2009;33(9):1648–61. http://dx.doi.org/10.1016/j.jedc.2008.12.007.
[13] Andersen KS, Termansen LB, Gargiulo M, Ó Gallachóir BP. Bridging the gap using energy services: Demonstrating a novel framework for soft linking top-down and bottom-up models. Energy 2019;169:277–93. http://dx.doi.org/10.1016/j.energy.2018.11.153.
[14] Anderson DG. Iterative procedures for nonlinear integral equations. J ACM 1965;12(4):547–60. http://dx.doi.org/10.1145/321296.321305.
[15] Gal T, Nedoma J. Multiparametric linear programming. Manag Sci 1972;18(7):406–22. http://dx.doi.org/10.1287/mnsc.18.7.406.
[16] Borrelli F, Bemporad A, Morari M. Geometric algorithm for multiparametric linear programming. J Optim Theory Appl 2003;118(3):515–40. http://dx.doi.org/10.1023/B:JOTA.0000004869.66331.5c.
[17] Fattahi A, Reynès F, van der Zwaan B, Sijm J, Faaij A. Soft-linking a national computable general equilibrium model (ThreeME) with a detailed energy system model (IESA-Opt). Energy Econ 2023;123:106750. http://dx.doi.org/10.1016/j.eneco.2023.106750.
[18] Ross AG, Connolly K, S. Vögele W Kuckshinrichs. A macro-level analysis of the socio-economic impacts of climate change driven water scarcity: Incorporating behavioural and resilience aspects. Water Res X 2024;23:100223. http://dx.doi.org/10.1016/j.wroa.2024.100223.
[19] McGregor PG, Ross AG, Swales JK. How fiscal policies affect energy systems: the importance of an 'environmental social wage'. Reg Stud 2021;55(8):1354–64. http://dx.doi.org/10.1080/00343404.2021.1893895.
[20] Loulou R, Lavigne D. MARKAL model with elastic demands: Application to greenhouse gas emission control. In: Carraro C, Haurie A, editors. Operations research and environmental management. Dordrecht: Springer; 1996, p. 201–20. http://dx.doi.org/10.1007/978-94-009-0129-2_9.
[21] Kehoe TJ. Multiplicity of equilibria and comparative statics. Q J Econ 1985;100(1):119–47. http://dx.doi.org/10.2307/1885738.
[22] Kehoe TJ, Whalley J. Uniqueness of equilibrium in large-scale numerical general equilibrium models. J Public Econ 1985;28(2):247–54. http://dx.doi.org/10.1016/0047-2727(85)90072-6.
[23] Brent RP. Algorithms for minimization without derivatives. Englewood Cliffs, NJ: Prentice-Hall; 1973.
[24] Deuflhard P. Newton methods for nonlinear problems: affine invariance and adaptive algorithms. Berlin, Heidelberg: Springer; 2011, http://dx.doi.org/10.1007/978-3-642-23899-4.
[25] Rump SM. Verification methods: Rigorous results using floating-point arithmetic. Acta Numer 2010;19:287–449. http://dx.doi.org/10.1017/S096249291000005X.
[26] Higham NJ. Accuracy and stability of numerical algorithms. second ed.. Philadelphia, PA: Society for Industrial and Applied Mathematics; 2002, http://dx.doi.org/10.1137/1.9780898718027.
[27] Lanz B, Rausch S. General equilibrium, electricity generation technologies and the cost of carbon abatement: A structural sensitivity analysis. Energy Econ 2011;33(5):1035–47. http://dx.doi.org/10.1016/j.eneco.2011.06.003.
[28] Scarf H. The approximation of fixed points of a continuous mapping. SIAM J Appl Math 1967;15(5):1328–43. http://dx.doi.org/10.1137/0115116.
[29] Ross AG. From transient shocks to unexpected outcomes: disruptive drivers in scenario pathways, preprint. 2026, http://dx.doi.org/10.48550/arXiv.2604.20879.