Critique of Engelhardt on planning

Allin Cottrell*

September 2, 2021

The specific contribution made by Engelhardt's paper[1] is a turbocharged version of the old 'too many equations to be solved' dismissal of central planning. Given the exponential growth in computing power since that argument was first made, in order to make his case he must come up with a vastly increased number of equations to be solved. In short, he does this by inventing an imaginary planning problem whose size is the product of the number of distinct goods and the number of consumers, then grossly overstating the complexity of this problem.

Here's a statement of the problem he sets (not a direct quotation but a good faith transcription into relatively formal language):

> The planners must allocate given total quantities (Q_j, $j = 1, \ldots, m$) of each of m goods across n consumers so as to maximize their total utility. The consumers' utility functions (assumed to be known by the planner) are quadratic and additive, so that the utility of consumer i is given by
>
> $$U_i = \sum_{j=1}^{m} \left(a_{ij} q_{ij} + b_{ij} q_{ij}^2 \right) \qquad a_{ij} \geq 0,\, b_{ij} \leq 0,$$
>
> where $q_{ij} \geq 0$ denotes the allocation of good j to consumer i. The problem is to maximize $\sum_i U_i$ by selection of $[q_{ij}]$ subject to $\sum_i q_{ij} = Q_j$ for all j.[2]

A few basic points must be made right away.

- The notion that the planners could know the utility functions of all consumers is absurd (as, indeed, is the notion that consumers themselves could know their own utility functions in full generality) and it is not a notion that has been entertained by any serious proponents of socialism.
- All historically existing socialist economies, and almost all socialist theorists, have seen the distribution of personal consumer goods as a job for a market of some sort (although the sphere of operation of this market is limited in relation to that in capitalist economies).
- The problem Engelhardt gives the planners is certainly not solved by the market. Consumers in the economics textbooks maximize their own utility (subject to their budget constraint) via their choice of which goods to consume in which quantities, but there's

*Department of Economics, Wake Forest University.

[1] Lucas Engelhardt, 'Central Planning's Computation Problem', *Quarterly Journal of Austrian Economics* 16:2, 2013, pages 227–246.

[2] What about the planning of production? Well, he's aiming for an argument *a fortiori*: if he can show that the pure allocation problem is insoluble, it's clear that the production-plus-allocation problem cannot be solved.

no market mechanism acting to maximize *total* utility. Given the inequality that's inherent in capitalist economies, getting anywhere near that optimum would require the very convenient assumption that the utility people derive from consuming a given quantity of any given commodity is proportional to their income.[3]

We could stop there, but it might be instructive to explore the actual complexity of the unreal problem set by Engelhardt, which is nothing like what he claims.

Engelhardt assumes that the planners would address his problem via Gaussian elimination, and he generously grants them access to the combined computing power of the TOP500 supercomputers as of 2013, on the order of 200 petaflops.[4] He then constructs three scenarios with different numbers of consumers and products and calculates the required compute time, as shown in the table below.

Scenario	Consumers	Goods	Compute time
1	1000	1000	3 seconds
2	3×10^8	100	2.6 million years
3	6×10^9	80000	10^{19} years

(As Engelhardt notes, 6×10^9 is roughly the world's population and 80,000 is roughly the number of items in the US Consumer Price Index.)

His arithmetic is OK, but his assumption that Gaussian elimination would be used is wrong. This method has a complexity of $O(N^3)$ where N indicates the number of variables, and in context N would equal the number of consumers mutiplied by the number of goods. Take the cube of a huge number of variables (4.8×10^{14} in Scenario 3) and you get something that will indeed take a very long time to calculate. But Gaussian elimination is not even applicable to Engelhardt's problem. It's an instance of quadratic programming (QP)—the maximization of a quadratic objective function under linear constraints. QP reduces to a linear problem only if the constraints are all equalities, but here we have nm inequality constraints (ruling out negative allocations to consumers).

In general the complexity of a QP problem can be difficult to assess (at worst such problems are NP-hard). But although this case has a huge number of constraints it also has a rather simple, regular structure which makes it amenable to solution by a relatively straightforward algorithm—based on the principle that we can maximize total utility by equating the marginal utility of each good across the consumers. Here is a suitable algorithm; the key ingredient is at step 4, where we iteratively reallocate goods away from those with below-average, and towards those with above-average, marginal utility, assessed at the current allocation.

1. Let $k = 0$. Allocate each consumer an equal share of each good: $q_{ij} = Q_j/n$ for all i, j. Calculate the baseline mean utility, $\bar{U}^0 = n^{-1} \sum_i U_i$.
2. If $k > 0$, recalculate mean utility, $\bar{U}^k = n^{-1} \sum_i U_i$, and if the increase over the previous iteration falls below some small threshold, stop.
3. For each good calculate the per-consumer marginal utilities, $MU_{ij} = a_{ij} + b_{ij} q_{ij}$, and the mean across consumers, $\langle MU_j \rangle = n^{-1} \sum_i MU_{ij}$.

[3] The most ambitious provable optimality result for a competitive market economy is that—under a set of stringent and unrealistic assumptions—it delivers a *Pareto-optimal* allocation. That means, an allocation such that you can't make anyone better off without making someone else worse off. Pareto optimality does not equal maximized utility. It's probably obvious to most people that making the super-rich a bit worse off to make a lot of poor people better off would increase total utility.

[4] Petaflop: a computing speed of 10^{15} floating-point operations per second.

4. Calculate the revised allocations, $\hat{q}_{ij}$, that would move all consumers to the mean marginal utility for each good, using the inverse of the MU functions: $\hat{q}_{ij} = (\langle MU_j \rangle - a_{ij})/2b_{ij}$. If any of the $\hat{q}_{ij}$ are negative, set them to zero. (A negative $\hat{q}_{ij}$ indicates that consumer i does not value good j highly enough to participate in the optimal allocation of that good.) By equalizing MU across participating consumers we are increasing "notional" utility, but the resulting allocations may not respect the adding-up constraints, so scale the allocations for each good uniformly: $q_{ij} = (Q_j / \sum_i \hat{q}_{ij})\, \hat{q}_{ij}$. Let $k = k + 1$ and go to step 2.

It's clear that the calculations involved at each iteration are $O(nm)$, i.e., linear in the total number of unknowns. The remaining factor is the number of iterations, k, needed to get close enough to the maximum. Given the quadratic objective function, this should be small.

I tested this algorithm on a basic desktop PC, using a single Intel i7 core. The utility coefficients were generated as random draws from normal distributions, with parameters such that (a) most consumers had positive but diminishing marginal utility over the relevant range of quantities and (b) no consumer had indefinitely increasing marginal utility. The run-times for various sizes of the problem were as follows: 0.6 seconds with $n = m = 1000$; 6 seconds with $n = 1000$ and $m = 10,000$; and about a minute with $n = m = 10,000$. These times confirm the linear time-order of the algorithm in practice. The required number of iterations was 3 or 4, and appears to be independent of the scale of the problem.

The largest problem I simulated is about 5 million times smaller than Engelhardt's monster Scenario 3 (6 billion consumers and 80,000 goods). But the assumed petaflop supercomputer has a speed on the order of a million times greater than a single i7 core. Given that my case ran in a minute, one would expect the monster case to run in about 5 minutes on a petaflop machine; compare Engelhardt's hyperinflated figure of 10^{19} years. And in this comparison I've ignored the fact that the supercomputers of 2021 are a few orders of magnitude faster than in 2013.

None of the above is intended to give credence to Engelhardt's formulation of the problem faced by planners. The object is simply to point out that his invented problem is not of anything like the complexity he claims. To put it bluntly, if you want to contribute to debate on this topic you have to know what you're talking about and it's clear that Engelhardt does not.

One more comment for anyone who's inclined to follow further. There's a footnote in Engelhardt's text (p. 231) where he seems to catch a glimmer of the algorithm described above:

> An alternative method is possible. The computer could begin with an arbitrary distribution of goods, and then consider possible trades and "swap" goods whenever a trade would be mutually beneficial. In order to be economically efficient, this routine would have to be computationally intensive, as the computer must consider a long chain of possible trades—the type of chain that, in a monetary economy, would be facilitated by the use of a medium of exchange. An interpersonal comparison of utility allows for a simpler algorithm: maximizing total social utility.

Here he seems to have forgotten his assumption that the planners know the consumers' utility functions. In the context of planning, no "long chain of possible trades" is required; simply, for each i, reallocate some of good i away from consumers for whom it has below-average marginal utility to those for whom it has above-average marginal utility. Such "trades" do *not* have to be mutually beneficial; they just have to increase total utility. In the market context no long chain of trades, however facilitated by a medium of exchange, will lead to maximum total utility, since trades will indeed be limited to the mutually beneficial subset.

Engelhardt has, commendably, given a hostage to fortune in allowing “interpersonal comparison of utility”. Such comparison licenses the thought: might it not increase total social utility to reallocate some resources from the super-rich to those struggling to make ends meet? (Something the market will never do.) But when he says that “maximizing total social utility” is an *algorithm*, I throw up my hands. It's an *objective*; show me your algorithm to achieve it! He doesn't specify an algorithm at all; he merely waves his hands and says it must be done via Gaussian elimination.