

When Does Saving Power Save the Planet?

Jackson Woodruff
J.C.Woodruff@sms.ed.ac.uk
University of Edinburgh

Michael F.P. O’Boyle
mob@inf.ed.ac.uk
University of Edinburgh

David Schall
david.schall@ed.ac.uk
University of Edinburgh

Christopher Woodruff
Christopher.woodruff@qeh.ox.ac.uk
University of Oxford

ABSTRACT

The computing industry accounts for 2% of the world’s emissions. Power-efficient computing is a frequent topic of research, but saving power does not always save the environment. Jevons’ paradox states that resource savings from increases in efficiency will be more than compensated for by increased demand by a process called *rebound* – making these ineffective ways to decrease emissions.

This is not the case for all applications within computing: applications whose demand is inelastic with respect to power consumption *can* have reduced power consumption. We analyze several large fields within computer science, including ML, the Internet and IoT, and provide directions on where power efficiency savings will help reduce carbon emissions.

We present the economic tools needed to decide whether power-efficiency improvements are likely to result in reduced or increased emissions. We conclude that many problems in computer science do have characteristics of rebound, meaning that green energy is the only solution for many fields.

ACM Reference Format:

Jackson Woodruff, David Schall, Michael F.P. O’Boyle, and Christopher Woodruff. 2023. When Does Saving Power Save the Planet?. In *2nd Workshop on Sustainable Computer Systems (HotCarbon ’23)*, July 9, 2023, Boston, MA, USA. ACM, New York, NY, USA, 6 pages. <https://doi.org/10.1145/3604930.3605719>

1 INTRODUCTION

The climate breakdown is increasingly causing heatwaves, droughts and floods and causing food and water insecurity for millions [6]. The computing-industry plays a non-insignificant part in this role, contributing 4% of the world’s electricity consumption and 2% of the world’s carbon emissions [23]. These emissions are forecast to continue increasing [2]: understanding the trends that define these emissions is critical to effectively combating them.

Power efficiency is considered a key metric of sustainability within datacenters [21], networks [32, 50], IoT [49] and other systems. Nearing 80 years of computing research, energy-efficiency gains have paradoxically resulted in *increased* power consumption and environmental impact and are forecast to continue to do so in many cases [48]. Numerous energy-saving strategies have been devised [10], yet to transform these into impact on carbon footprint,

they need to be targeted appropriately; without understanding when these techniques *actually* help the environment, they come across as greenwashing at best [3–5, 39].

These energy-saving techniques involve process-scheduling (e.g. [22]), more efficient use of hardware (e.g. [24]) and other system improvements (e.g. [38]). While these energy-saving techniques have been effective in reducing costs, they have largely not been effective in reducing the carbon footprint of computing. The rebound effect where increased demand offsets reduced prices is responsible for this, and overcoming it is one of the five biggest challenges in green computing [30].

This paper covers the economic framework used to frame the problem of rebound, which provides a basis to classify problems as *elastic* and *inelastic* and enables determination of direct rebound effects. We cover the limitations of this approach, and discuss indirect rebound effects and how to reason about those. We reach the following conclusions:

- Direct rebound is present in many of the most important applications in computing, but may be mitigated in some cases by monopolies.
- Using green energy is critical.
- Computing has potential for significant positive indirect rebound effects due to the replacement of more carbon-intensive activities.

2 DEFINITIONS

Global warming is truly a political problem. This does not mean that technology improvements do not play a role, but technology improvements often do not result in a reduction in emissions due to market behaviours. Economics is a critical tool to understand the impact of technology improvements on carbon emissions. There are several key terms to define in the context of computing:

Price Price refers to the cost of purchasing some product. We refer to price as P and change in price as ΔP .

Demand Demand refers to number of purchasers for a particular product (e.g. a laptop), or compute power (e.g. in AWS). We refer to demand as Q below and change in demand as ΔQ .

Rebound If some computation task is made more efficient, the energy required to perform it decreases, which reduces the price by some ΔP . However, the act of it becoming cheaper makes it useful in more cases, increasing the number of uses by some ΔQ . As per [12], we have a rebound of $Z\%$ if $Z\%$ of the savings from the energy-saving technique are consumed by increased demand. 100% rebound means that the efficiency improvement had no effect, and greater than

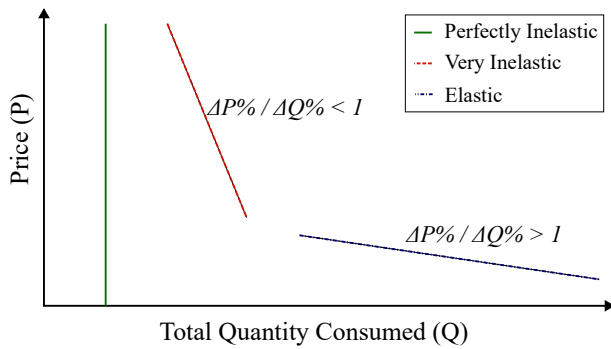


Figure 1: How products with different elasticities respond differently as costs of running change. $\Delta P\%$ and $\Delta Q\%$ refer to the percentage change in price and percentage change in demand respectively. The coefficient of elasticity is the ratio between these two, defined as: $\epsilon = (\Delta Q/Q)/(\Delta P/P)$. When we have inelastic behaviour, the coefficient of elasticity is less than one, and when we have elastic behaviour the coefficient of elasticity is greater than one.

that means the efficiency improvement had negative impact on total energy consumption.

Coefficient of Elasticity The coefficient of elasticity is a way of measuring how demand changes with price. It is defined as $\epsilon = (\Delta Q/Q)/(\Delta P/P)$ where P is price and Q is the quantity demanded.

Elastic Demand Elastic demand refers to demand that is highly sensitive to price changes: if a product's price reduces, the demand for that product will increase proportionately more than the price. In these cases, we have $\epsilon > 1$. The "Elastic" line on figure 1 shows how price and quantity consumed vary in elastic goods: we can see that a small drop in price produces a large increase in total quantity consumed.

Inelastic Demand Inelastic demand refers to demand that is insensitive to price changes: if a product's price reduces, the demand for that product will increase proportionately less than the price. In these cases we have $\epsilon < 1$. The "Very Inelastic" line on figure 1 shows how price and quantity consumed vary in inelastic goods: we can see that a large drop in price produces a modest increase in total quantity consumed.

Utility Utility refers to the "usefulness" of a product in an abstract manner. The usefulness may be different for different parties, but largely, increasing utility without increasing price will mean more demand.

Figure 2 shows an example of rebound – despite power-efficiency of individual transistors increasing by a factor of approximately 16 over the period of 2000–2005, the total power consumption of servers increased by a factor of 260%. Of course, not all of the growth was driven by increased power-efficiency of transistors, but Dennard scaling and Moore's law have underpinned a large part of the historic growth of the computer industry.




Figure 2: The power consumption of servers worldwide [37] plotted against Dennard scaling in the same window. The energy-efficiency improvements from the Dennard scaling have rebounded in the server market over the period of 2000–2005, with an increase in power consumption of servers of 2.6x, giving us a rebound of 260%. The effects of Dennard scaling are estimated using TSMC's leading-edge process nodes throughout the period of 2000–2005 and are normalized to 2000 levels.

3 DEMAND IN COMPUTING

Demand is typically measured with respect to cost. Reducing power consumption has an effect on cost, but it does not completely control the cost. This section explores how demand changes with in cases where power-consumption is not a significant cost, and also in cases where reductions in cost might not be passed on to the consumer.

3.1 Monopolies: Reductions in Cost vs Reductions in Price

In monopolistic markets, power-saving techniques may not result in a cost reduction for the end-user, meaning that the effect on the demand that we have discussed may not occur. The development of a power-saving technique that reduces the cost of providing a service *does not* need to be passed onto the consumer as there is no market to force the price down. In this case, monopolies may keep the prices they charge the same, despite their expenditures reducing.

In a non-monopolistic market, if a power/cost-saving technique is developed, the market will force the competing companies to lower their prices to account for these reduced costs of providing the service. This provides the basis for rebound: the lower prices enable consumers to consume more to make up for the reduction in prices.

Due to the low-cost to entry of software, true monopolies are rare in computing. However, near-monopolies are exceedingly frequent [26]. This is apparent with recent start-ups, where huge amounts of funding has been made available with the only aim being to grow user-bases and to make money from them later. These industries provide good opportunities to reduce carbon footprints, as they are unlikely to lose users to competitors, and will be eager to reduce costs without needing to pass those onto consumers.

	Demand Type	Power-Independent?	Power Savings -> Reduced Carbon Footprint?
A	Elastic Demand	No	No
B	Elastic Demand	Yes	Yes
C	Inelastic Demand	No	Yes
D	Inelastic Demand	Yes	Yes

Table 1: A table for classifying whether power-saving techniques will result in reduced energy use. The extent to which power-savings result in emissions reductions depends on the coefficient of elasticity and is covered in section 5.1.

3.2 Power-Independent Products

Typical definitions of rebound work with price changes, and in many products, power does not make up a significant fraction of the cost. We term these products *Power-Independent Products* which have the following properties:

- (1) Utility of the product is not limited by compute power.
- (2) Power is not a significant cost.

3.2.1 Elastic Demand. Typically, elastic demand would mean that power-saving techniques are ineffective at reducing carbon footprints. However, if a product has power-independent performance, power-saving techniques have no impact on the consumer. In these cases, even though a product can be elastic, it will not be elastic with respect to power consumption, making power-saving techniques effective.

3.2.2 Motivating Change with Power-Independent Products. Motivating change in these fields is challenging: those responsible for power-efficiency in such cases have little-to-no incentive to make their devices more power-efficient. This is not to say that the costs are irrelevant to the environment [49].

In other areas, green certification has been an important way to address this problem. The Energy Star program (discussed below) has worked well for home-related energy efficiencies. Similar, *trustworthy* green certifications are needed for fields where power consumption is a low cost.

There are a number of existing schemes. The Energy STAR program was a consumer-visible label that indicates a threshold of energy-efficiency, estimated to reduce the EU's power consumption by 0.4% [2]. The BlueAngel scheme targeted at datacenters is similar, although it is significantly less-effective [2, 31] in part due to poor choice of metric and in part due to the costs of installing measurement equipment. Significant lessons on the design of such schemes are presented by the European Commission in [2].

3.2.3 Taxonomy of Elasticities for Power Consumption. Table 1 shows how different types of application respond to power-saving techniques. Inelastic applications are the best target for power-saving techniques, along with power-independent products. The magnitude to which carbon savings are possible depends on the value of the coefficient of elasticity, and equations to calculate this are provided in section 5.1.

4 CASE STUDIES

In this section we will review several different case studies. We should note that predicting elasticity into the future of a particular problem is, at best, an informed guess. Instead of promising that these applications will always stay elastic/inelastic as described, we will look at current predictions where they are available and recent data where they are not.

For any of the below examples, new innovations (or stalled innovations) may change how demand may shift with increased performance.

4.1 A: A Case Study of Elastic Demand: Machine Learning

Machine learning models have exploded in size, with the largest models costing millions of dollars to train [27], much of this spent on power consumption. Clearly, the demand for more efficient ML hardware is there [18], but as training becomes more efficient, the models trained have *wider applicability*, meaning they are more valuable, increasing the value of training them, and therefore the money that will be spent on it and the carbon footprint.

For example, Google reports a 2.7x performance improvement per Watt of their latest generation TPU v4 over its predecessor TPU v3 [35]. At the same time, OpenAI increased the number of parameters for its latest GPT-4 model by one order of magnitude [9] – an indication that these power-efficiency improvements are not resulting in power-savings.

4.2 A: A Case Study of Elastic Demand: Bitcoin

Bitcoin is a classic elastic computing application [33]. Over 80% of the cost of bitcoin is the energy consumption [47], and there is an obvious, direct benefit to making mining cheaper: profit margins become bigger and more money can be spent on growing the mining operation. Any power-saving technique will simply result in a larger ASIC-cloud [36] for bitcoin mining.

In these cases, there are only two ways forward. The first is government regulation to increase the cost of energy to reflect the cost to society. It is important to note that under these circumstances, the rebound effect is increased [14]¹. The second is to use renewable energy. Already, approximately 30% of bitcoin's energy comes from renewable sources [47].

However, in summary, for truly elastic applications, the only ways forward are legislation to reduce their negative impacts on the environment or to encourage more use of renewable energy.

4.3 C: A Case Study of Inelastic Demand

Business compute is a good example of inelastic demand. For many businesses, compute performs required tasks (e.g. providing all employees with laptops), and the performance and cost of the individual computers is small with respect to the overall operating costs. Further, provisioning of such compute is often complex and operates on fixed schedules (e.g. replacing all laptops every X years).

Various works have inspected elasticities in business computing. Jiang et al. [34] conclude that purchases of computers by businesses

¹The overall price might be increased by increases in the cost of energy, but in a case where the cost is already increased, an energy-saving technique reduces the price more, resulting in more rebound.

are very inelastic (with a coefficient of 0.51)² and other studies have reached similar conclusions [45]. Analysis of consumer purchases of laptops shows that elasticity falls with laptop age [43] (the logic being that new laptops are luxury items, while those purchasing laptops that are no longer cutting-edge are more likely to do so out of necessity).

4.3.1 Caveats of Inelastic Demand. Particularly in computing, inelastic demand evolves through time with innovations. The core challenge to overcome here is that a field that was previously inelastic may *become* elastic with a new innovation. Even in the business example inspected above, estimates of inelasticity either require short timeframes (as in [34]) or account for the complex patterns of growth computing usage (as in [45]).

4.4 B: A Case Study in Elastic Power-Independent Products

The Internet is a classic study of rebound: traffic has traditionally grown with capacity [17] and internet-access plans are highly elastic [11]. However, the costs of running the internet are not dominated by power consumption as much as they are by the challenges of maintaining a large network that is robust to traffic bursts, buggy components and natural disasters. An example of this is in Telecom Italia, which was reported in 2018 to need more than 2TWh of electricity to run, nearly: 1% of Italy's national energy demand [13]. Given the costs of electricity in Italy, and the reported turnover of Telecom Italia, this is approximately 5% of expenditures, far from dominating the costs despite playing a non-insignificant role on a national scale.

The result is that although power-efficiency is a relevant metric to the capacity of the Internet, it is not a defining metric. An example of this can be in AT&T, one of the largest telcos in the world. Over the time-period of 2018–2022, AT&T's power consumption decreased by 14% [7], despite an increase in traffic by a factor of two [8].

This situation corresponds well to a classic example of inelasticity: the effect of fuel-efficiency in cars in enabling more driving. The effect of increased fuel-efficiency has not been a dramatic increase in driving [42] as other factors, such as the time-cost of driving are much more limiting [44].

4.5 D: A Case Study in Inelastic Power-Independent Products

The Internet-of-Things (IoT) is an excellent example of a field where power consumption is not a *relevant* cost to the producer. IoT devices are sold to consumers, who pay for the electricity costs, but while this energy consumption adds up on a larger scale [49], for each individual consumer the power consumption is not likely to be very significant.

Power-saving techniques directed at IoT devices are likely to be highly effective, as consumers are unlikely to purchase more IoT devices due to lower power consumption (e.g. smart doorbells).

²Jiang et al. [34] do conclude that cloud computing is moderately elastic – a conclusion backed up by studies on the rebound of cloud computing environments [41], which has significantly reduced costs over traditional enterprise computing [1].

Note that this differs from a business use case: businesses tend to be much better at accounting for the lifetime costs of a product.

Again, this case of inelasticity has a corresponding area within economics. Studies of efficient LED lightbulbs show very little rebound effect within consumer behaviour [40].

5 TOOLS TO IDENTIFY THE ELASTICITY OF PROBLEMS

Inelastic problems are characterized by several key features. For example, a third factor can dominate the uptake rate of a product. An example of this is speed cameras, whose installation is typically limited by politics rather than the direct cost of installing or running the cameras.

The largest category of inelastic goods are those that are considered necessities. Mobile phones are increasingly becoming necessities, and so are inelastic in the sense that making them cheaper will not significantly increase the number purchased (in rich countries).

5.1 Empirically Determining Rebound

$$\epsilon = \frac{\Delta Q/Q}{\Delta P/P} \quad (1)$$

Where Q is the demand, P is the price and $\Delta Q, \Delta P$ are their changes respectively. This equation holds for either linear price/demand or for small ΔP . Using this equation, we calculate the coefficient of elasticity, ϵ . $\epsilon < 1$ means that the product is inelastic, and $\epsilon > 1$ means that the product is elastic and so susceptible to rebound.

5.1.1 When Does Saving Power Reduce Power Demand. Given a product with a particular coefficient of elasticity ϵ_p , and an energy-efficiency technique that reduces the power consumption of this product by a factor of $e \leq 1.0$ so that $e = 1$ means that power consumption has been eliminated, we can analytically decide whether this innovation will result in overall reduced carbon emissions. We also define I , the fraction of the cost that comes from power-consumption as:

$$I = \frac{\psi_P}{\psi_P + \phi_P}$$

Where ψ_P is the cost of power and ϕ_P is the non-power cost.

ΔP and e are related by the equation:

$$\Delta P = e \times P \times I$$

Using equation 1 we can write:

$$\begin{aligned} \Delta Q &= \epsilon_p (\Delta P/P) Q \\ &= \epsilon_p \times e \times I \times Q \end{aligned} \quad (2)$$

To compute the environmental impact of this change, we should consider the embodied emissions, ϕ_E and the power-consumption emissions ψ_E per unit of consumption of the original product. The change in emissions is given by:

$$\Delta Q(\phi + (1 - e)\psi) - Q\psi_e$$

Which we can rewrite using equation 2 to get a change in emissions of:

$$\underbrace{\epsilon_p \times e \times I \times Q(\phi_E + (1 - e)\psi_E)}_{\text{Rebound}} - \underbrace{Q\psi_{Ee}}_{\text{Direct Improvement}}$$

5.1.2 Non-Linear Demand. For large price changes, linear demand often does not make sense. The definition of elasticity in equation 1 still holds, but ϵ , the elasticity, will vary with ΔP . This means that we must find other ways of estimating ϵ when the price change is significantly different than what we have data for. In general, we can treat ϵ as a function of the new and old prices:

$$\underbrace{\epsilon_P(P, P - \Delta P) \times e \times I \times Q(\phi_E + (1 - e)\psi_E)}_{\text{Rebound}} - \underbrace{Q\psi_E e}_{\text{Direct Improvement}}$$

5.2 The Limits of Projecting Elasticities

There are several key limits to understand when projecting elasticities using equations. First, when price changes are large, the potential for a particular technology to be used changes significantly. This is a reflection on the gradient of the price/demand graph being taken at a single point for these equations.

As prices change significantly, gradients on the price/demand graph are likely to change significantly as entire new usecases are opened and demand for existing usecases is saturated. Despite these limitations, the equations presented above can provide an educated guess as to what the rebound will be.

6 INDIRECT REBOUND

The wider challenges of addressing this consumption-based environment is one of the biggest challenges in addressing climate change, and is not the focus of this article, which seeks to explain how we can reduce the footprint of the *computing* industry in particular, and when we might be inadvertently increasing that footprint.

The indirect rebound effect is where money that is being saved on computing due to a decrease in price is spent on other, emitting, items. For example, If a typical home saves a small amount of money every year on power consumption due to more efficient IoT devices, that money will be spent somewhere else: perhaps on flying. This indirect rebound effect has been estimated at approximately 10–15% [16].

The other key element of indirect rebound effect is the *substitution effect*, where price changes drive consumers to substitute one good with another. We will discuss this further in section 6.3.

6.1 Rebound in the Long Term

The rebound effect also refers to long-term effects of wider innovations. However, many long-term innovations are impossible to predict (almost by definition of innovation). For example, the coal that Jevons saw rebound effects in was replaced by gas, petrol and diesel with various further innovations.

The same applies in computer science: while we can reliably assume that machine-learning will continue to use every watt available to it, systematic reductions in the per-use cost of IoT that do result in reduced power consumption may suddenly open an entirely new usecase that we cannot (yet) foresee.

6.2 Is Rebound Overrated?

We do *not* set out to argue that energy-efficiency improvements are bad. There are many benefits to increasing power-efficiency, from increased computing power to solving new problems — including

using computers to reduce the carbon footprint of other tasks [20]. In this sense, rebound is overrated: in many cases, the indirect effects are more important than the direct power-consumption effects.

6.3 The substitution effect: Is Rebound Good?

There is a strong argument to be made that computing is one of the greenest industries. Money spent in virtual spaces (e.g. on video games) has virtually zero physical resource usage, especially when compared to alternative spaces where individuals spend their money.

The *substitution effect* plays a key role in this: when power-efficiency developments enable computers to take over a new chunk of consumer spending that was previously being spent on more polluting industries, the effects are good.

Further, although increasing power consumption may well be a problem, computers have a clear path to being turned green. Because we are considering an industry largely driven by electricity, the transition to renewable energy is relatively simple.

6.4 Embodied Carbon and Related Emissions

For many devices, particularly those with large batteries or screens that are not always on (e.g. phones, personal computers), embodied carbon is more important than the power usage while running the device [15, 19, 46].

The same lessons we have applied here can be applied to these cases, although some nuances apply:

- A frequent business model is one of ethical production, where *increased* costs are charged with consumers being promised devices that have lower environmental footprints³.
- Many of the environmental impacts embodied into devices are not just the carbon, but other precious metals such as gold and copper [29].

6.4.1 Disposal Charges. Computers contain rare metals, some of which are worth recycling [28]. However, many of the components within typical computer design are not worth the cost of recycling. Disposal charges have been introduced for other items such as bottles and have been used to set recycling targets for computers [25].

These charges have a similar impact to a carbon tax: they increase the cost of doing things that are damaging to the environment. The result is that techniques that reduce the carbon footprint are even *more* susceptible to rebound.

7 CONCLUSION

Understanding the potential for rebound is critical in understanding the potential of power saving techniques to save the planet. In this paper, we present economic models for understanding rebound effect, and present argument that in many sub-fields of computer science, rebound is unavoidable. We come to two key conclusions: that green energy is critical in reducing the carbon footprint of the computer industry, and that ensuring that when power-saving techniques are developed, that they are applied to fields where they may not directly be profitable.

³e.g. fairphone.com

The unfortunate reality for computing is that energy-efficiency gains are likely to result in rebound for many fields, and the growing nature of the discipline means that fields that do not show rebound are unlikely to constitute a significant part of computing's power consumption for very long.

On the positive side, sustainable electricity generation promises to solve many of the emissions problems in computing, and in this regard, some of the rebound face in computing is *good*. Digital markets are detached from physical resources in a way that makes it possible to achieve growth without many of the environmental concerns that come from physical products.

REFERENCES

- [1] Carbon disclosure project study 2011: Cloud computing — the IT solution for the 21st century. Available at https://www.att.com/Common/about_us/files/pdf/cdp_us_cloud_computing.pdf.
- [2] Energy-efficient cloud computing technologies and policies for an eco-friendly cloud market. Available at <https://digital-strategy.ec.europa.eu/en/library/energy-efficient-cloud-computing-technologies-and-policies-eco-friendly-cloud-market>.
- [3] 2021 nvidia corporate social responsibility report. Available at <https://images.nvidia.com/aem-dam/Solutions/documents/FY2021-NVIDIA-CSR-Social-Responsibility.pdf>.
- [4] AMD announces ambitious goal to increase energy efficiency of processors running AI training and high performance computing applications 30x by 2025. Available at <https://www.amd.com/en/press-releases/2021-09-29-amd-announces-ambitious-goal-to-increase-energy-efficiency-processors>.
- [5] Intel: Corporate responsibility report (product energy efficiency).
- [6] Climate change: a threat to human wellbeing and health of the planet. taking action now can secure our future. Available at https://www.ipcc.ch/report/ar6/wg2/downloads/press/IPCC_AR6_WGII_PressRelease-English.pdf.
- [7] Energy management data. Available at <https://sustainability.att.com/priority-topics/energy-management>.
- [8] Network quality and reliability data. Available at <https://sustainability.att.com/priority-topics/network-quality-reliability>.
- [9] The ultimate guide to gpt-4 parameters: Everything you need to know about nlp's game-changer. Available at <https://medium.com/@mlubbad/the-ultimate-guide-to-gpt-4-parameters-everything-you-need-to-know-about-nlps-game-changer-109b8767855a>.
- [10] ANAND, V., XIE, Z., STOLET, M., VITI, R. D., DAVIDSON, T., KARIMPOUR, R., ALZAYAT, S., AND MACE, J. The odd one out: Energy is not like other metrics. *Hot Carbon* (2022).
- [11] ATSUSHI, I. Estimating demand for cellular phone services in Japan. *Telecommunications Policy* (2005).
- [12] BERKHOUT, P. H. G., MUSKENS, J. C., AND VELTHUIJSEN, J. W. Defining the rebound effect. *Energy Policy* (2000), 425–432.
- [13] BIANCO, C., CUCCHIETTI, F., AND GRIFFA, G. Energy consumption trends in the next generation access network — a telco perspective. *International Telecommunications Energy Conference* (2007).
- [14] BRENNAN, T. J. Energy efficient policy puzzles. *The Energy Journal* 34 (2013).
- [15] BRUNVAND, E., KLINE, D., AND JONES, A. K. Dark silicon considered harmful: A case for truly green computing. *International Green and Sustainable Computing Conference* (2018).
- [16] CHITNIS, M., SORRELL, S., DRUCKMAN, A., FIRTH, S. K., AND JACKSON, T. Turning lights into flights: Estimating direct and indirect rebound effects for UK households. *EP* (2013).
- [17] COFFMAN, K. G., AND ODLYZKO, A. M. Internet growth: Is there a "moore's law" for data traffic? *Handbook of massive data sets* (2002), 47–93.
- [18] DALLY, W. J., AND KIRK, D. B. Evolution of the graphics processing unit (GPU). *Micro* (2021).
- [19] DENG, L., BABBITT, C. W., AND WILLIAMS, E. D. Economic-balance hybrid LCA extended with uncertainty analysis: case study of a laptop computer. *Journal of Cleaner Production* (2011).
- [20] DONTI, P. How machine learning can help tackle climate change. *XRDS* (2020).
- [21] GANDHI, A., GHOSE, K., GOPALAN, K., HUSSAIN, S. R., LEE, D., LIU, D., LIU, Z., MCDANIEL, P., MU, S., AND ZADOK, E. Metrics for sustainability in data centers. *Hot Carbon* (2022).
- [22] GAO, K., HUANG, Y., AADOLLAH, A., AND WANG, L. A review of energy-efficient scheduling in intelligent production systems.
- [23] GELENBE, E., AND CASEAU, Y. The impact of information technology on energy consumption and carbon emissions. *Ubiquity* (2015).
- [24] HACKENBERG, D., SCHONE, R., ILSCHKE, T., MOLKA, D., SCHUCHART, J., AND GEYER, R. An energy efficiency feature survey of the intel haswell processor. *IEEE International Parallel and Distributed Processing Symposium Workshop* (2015).
- [25] HANSELMAN, S. E., AND PEGAH, M. The wild wild waste: e-waste. *SIGUCCS* (2007).
- [26] HAUCAP, J., AND HEIMESHOFF, U. Google, Facebook, Amazon, eBay: Is the internet driving competition or market monopolization? *International Economics and Economic Policy* (2014).
- [27] HELLEDOORN, V. J., AND SAWANT, A. A. The growing cost of deep learning for source code. *Communications of the ACM* 65 (2022), 31–33.
- [28] HILTY, L. Electronic waste—an emerging risk? *Environmental Impact Assessment Review* 25 (2005), 431–435.
- [29] HILTY, L. M., KOHLER, A., SCHEELE, F. V., ZAH, R., AND RUDDY, T. Rebound effects of progress in information technology. *Poiesis Prax* 4 (2006), 19–38.
- [30] HILTY, L. M., AND LOHMANN, W. The five most neglected issues in "green IT". *Cepis Upgrade XII*, 4 (October 2011).
- [31] HORNER, N. C. *Powering the Information Age: Metrics, Social Cost Optimization Strategies, and Indirect Effects Related to Data Center Energy Use*. PhD thesis, Carnegie Mellon University, 2016.
- [32] JACOB, R., AND VANBEVER, L. The internet of tomorrow must sleep more and grow old. *Hot Carbon* (2022).
- [33] JALAN, A., MATROVSKYY, R., AND URQUHART, A. Demand elasticities of bitcoin and ethereum. *Economics Letters* (2022).
- [34] JIANG, Q., AND LEE, YOUNG CHOONA ND ZOMAYA, A. Y. Price elasticity in the enterprise computing resource market. *Cloud Computing* (2016).
- [35] JOUPI, N. P., KURIAN, G., LI, S., MA, P., NAGARAJAN, R., NAI, L., PATIL, N., SUBRAMANIAN, S., SWING, A., TOWLES, B., YOUNG, C., ZHOU, X., ZHOU, Z., AND PATTERSON, D. Tpu v4: An optically reconfigurable supercomputer for machine learning with hardware support for embeddings, 2023.
- [36] KHAZRAEE, M., MAGAKI, I., VEGA GUTIERREZ, L., AND TAYLOR, M. ASIC clouds: Specializing the datacenter. *IEEE Micro* (2017), 1–1.
- [37] KOOMEY, J. G. Estimating total power consumption by servers in the u.s. and the world. Tech. rep., Lawrence Berkeley National Laboratory and Stanford University, 2007.
- [38] PRIYA, B., PILLI, E. S., AND JOSHI, R. A survey on energy and power consumption model for Greener Cloud. *International Advance Computing Conference* (2013).
- [39] SANDS, H., AND MORISON, B. Greenwashing in the information industry. *The iJournal* 5 (2020).
- [40] SCHLEICH, J., MILLS, B., AND DUTSCHKE, E. A brighter future? quantifying the rebound effect in energy efficient lighting. *Energy Policy* 72 (2014), 35–42.
- [41] SEDLACKO, M., MARTINUZZI, A., AND DOBERNIG, K. A systems thinking view on cloud computing and energy consumption. *2nd International Conference on ICT for Sustainability* (2014).
- [42] SMALL, K. A., AND VAN DENDER, K. Fuel efficiency and motor vehicle travel: The declining rebound effect. *The Energy Journal* 28 (2007), 25–51.
- [43] STAVINS, J. Estimating demand elasticities in a differentiated product industry: The personal computer market. *Journal of Economics and Business* (1997).
- [44] TAIEBAT, M., STOLPER, S., AND XU, M. Forecasting the impact of connected and automated vehicles on energy use: A microeconomic study of induced travel and energy rebound. *Applied Energy Journal* 247 (2019), 297–308.
- [45] TAM, K. Y., AND HUI, K. L. Price elasticity and the growth of computer spending. *Transactions on Engineering Management* 46 (1999), 190–200.
- [46] TEEHAN, P., AND KANDLIKAR, M. Comparing embodied greenhouse gas emissions of modern computing and electronics products. *Environmental Science and Technology* 47 (2013), 3997–4003.
- [47] TRUBY, J., BROWN, R. D., DAHDAL, A., AND IBRAHIM, I. Blockchain, climate damage, and death: Policy interventions for reduce the carbon emissions, mortality, and net-zero implications on non-fungible tokens and bitcoin. *Energy Research and Social Science* 88 (2022).
- [48] VILLA, O., JOHNSON, D. R., O'CONNOR, M., BOLOTIN, E., NELLANS, D., LUITJENS, J., SAKHARNYKH, N., WANG, P., MICIKEVICIUS, P., SCUDIERO, A., KECKLER, S. W., AND DALLY, W. J. Scaling the power wall: A path to exascale. *International Conference for High Performance Computing, Networking, Storage and Analysis* (2014).
- [49] WANG, W., SOBRAL, V. A. L., BILLAH, M. F. R. M., SAODA, N., NASIR, N., AND CAMPBELL, B. Low power but high energy: The looming costs of billions of smart devices. *Hot Carbon* (2022).
- [50] ZILBERMAN, N., SCHOOLER, E. M., CUMMINGS, U., MANOHAR, R., NAFUS, D., SOULE, R., AND TAYLOR, R. Toward carbon-aware networking. *Hot Carbon* (2022).