https://pubsonline.informs.org/journal/ited

INFORMS TRANSACTIONS ON EDUCATION

Vol. 00, No. 0, Xxxx 0000, pp. 000–000

Submitted to *INFORMS Transactions on Education*

ISSN 1532-0545, EISSN 1532-0545

Reverse perspective: On subjectivity in mathematics education

Christopher Thomas Ryan
UBC Sauder School of Business, University of British Columbia, chris.ryan@sauder.ubc.ca

Authors are encouraged to submit new papers to INFORMS journals by means of a style file template, which includes the journal title. However, use of a template does not certify that the paper has been accepted for publication in the named journal. INFORMS journal templates are for the exclusive purpose of submitting to an INFORMS journal and are not intended to be a true representation of the article's final published form. Use of this template to distribute papers in print or online or to submit papers to another non-INFORM publication is prohibited.

Abstract. Mathematics is one of the fundamental building blocks of modern OR/MS. This makes the teaching of mathematical concepts a central challenge for our field. I argue that a standard approach to mathematics education—thinking about math as a powerful tool for coming to a "correct" solution—has major limitations in our AI-fueled era. At the introductory, intermediate, and some might even say advanced, levels, essentially every problem presented to students whose primary challenge is uncovering a "correct" solution or proof can easily be handled with the use of AI tools. One response has been to go back to a "pencil and paper" era of supervised mathematical performance. I propose an alternate direction, the "reverse perspective." Instead of thinking of mathematics as a tool for finding "correct" solutions and proofs, it can be seen as a creative language to structure and think about the world from each student's own subjectivity. I provide a philosophical defense of this perspective grounded in my own personal teaching experiences and provide concrete examples of teaching from the "reverse perspective".

Funding: This research was supported by the Social Sciences and Humanities Research Council of Canada Insight Development Grant.

Key words: mathematics education, subjectivity, mathematical creativity

Pre-amble

This is an essay about teaching mathematics. The primary example I will discuss is a basic example of multiplication: "$4 \cdot 2$". This may seem a strange thing to bring to a journal devoted to teaching operations research, management science, and operations management. In this preamble, I will say a few words about why I think it belongs here.

When I looked closely at what I do in my own classroom, a not-insignificant portion of it is teaching elementary mathematical concepts. It is not mathematics as a mathematics department would recognize it (there are no proofs, and even the hardest calculations can easily be automated), but it is mathematics all the same.

I am teaching this content, for the most part, to people who do not come to business school to learn mathematics. My MBA and undergraduate business students are bright and capable, but many are not what we might call "math people." The last math course they took was likely in high school or first year of college, and they probably stopped thinking seriously about the meaning of math long before that. I suspect most readers of this journal know these students well.

Our field has known for a long time that formulas and "word problems" are not the hardest thing we teach. Morris (1967) wrote of the "art" of modeling almost sixty years ago, and Powell (1995, 1998) built a course around teaching that art the way painting is taught, by doing it under the eye of someone more experienced. Grossman (2001), writing in the second issue of this journal, argued that the business school management science course was in decline in part because it had been organized around algorithms rather than around the ill-defined problems managers actually face.

Each of these authors, in their own way, describes a tension between mathematics as a set of procedures that converge on an answer and mathematics as a way of looking at a messy situation. In the second, there is no single answer, and the quality of a model is something about which reasonable people might disagree. This essay provides a metaphor for this tension, borrowed from art history, and asks what might happen if we leaned deliberately into the "mess" even when covering the "basics".

There is also a more pressing motivation for writing this essay. Many of us are rethinking our courses now that AI tools can formulate, solve, and code a great many of the problems we used to assign. Recent papers in this journal describe some of these redesigns (Zhang 2026, Daw et al. 2026). I contend that one possible change is to reverse the direction in which we ask students to inquire, moving them away from the answer and toward the world.

A word on how to read what follows. The essay is written for a general reader, and it uses the simplest mathematics I could find, the multiplication of whole numbers, so that anyone can try the exercises without preparation. I have offered this framing and discussion many times in business school classrooms. I hope readers of this journal will forgive the elementary setting and, as they go, substitute their own favorite topics: linear programming, Little's Law, the newsvendor model. The appendix offers some thoughts that might assist in this substitution.

* * *

* * *

My life as a teacher of mathematics

I have been an informal teacher of mathematics since I was a child. I would get up from my desk and help other kids during math class, to the point it sometimes got me in trouble. During high school, to the occasional chagrin of my mother, I would frequently have classmates over to study for exams. Even now, I teach math to my two sons (and some of their friends) in two weekly classes.

I believe the subject of mathematics is deep and important. But I also grew to greatly dislike mathematics at times, largely because of how it is taught and viewed as an intellectual subject. I am writing a whole book about this, but the essential argument is that my predominant experience of mathematics growing up was of it being used as a kind of intellectual measuring stick, one that rewards the "smart" and chastens the "weak". I wavered between these two, being "good" enough to appear "smart" while still feeling the sting of punishment in myself and, more empathetically, through others.

But now AI has come along and put many thing about mathematics in question. I have read already at least ten posts on social media questioning the future of mathematics research and education. Is it possible that mathematics has somehow gone from one of the hardest things to do to one of the easiest, with the "proof and problem-solving calculator" we call ChatGPT Astra or Claude Fable (insert the latest model whenever you read this)? This essay explores some of the implications of this change, and why I think we might all be better off for it. In particular, I am hopeful that this change will give rise to a democratization of the depth and beauty of mathematics.

Math as measuring stick

Mathematics is arguably the most important subject taught in school. I do not necessarily mean that it is the most useful subject, although it might be. I mean that it may be the most consequential for opening up—and closing off—future paths. Mathematics is unusually cumulative: doing well in one course gives access to the next, while falling behind can narrow the courses, programs, and eventually careers available to a student. Every "selective" school I know of places a strong emphasis on mathematics in its entrance requirements. Mathematics is also unusually competitive as a school subject, with a whole ecosystem of contests, enrichment programs, rankings, and even the concept of the "mathlete". In this sense, mathematics is something students learn in school, but it is also one of the principal ways that schools sort them. It is important to understand why.

First, mathematics benefits from "objectivity". If you ask people about a defining feature of math, a common answer is that "there is one answer." Theorems are either true or false.[1] The truthness

[1] I know this is wrong; some questions are undecidable, but we digress.

or falseness of a statement is not a matter of perception. It simply "is" and you (the learner) must adjust. When I left high school, before going to my undergraduate studies and intending never to study math again, my math teacher said I would miss the "clarity" of mathematics, referring to its objectivity and determinism. I didn't.

But this objectivity also has a societal "benefit". When there is a single answer, it is easy to grade people's performance. You can give students a long list of questions, all easy to mark—just see whether they got the final number—and see who answers more correctly. Those are your good students. There is no credible argument from the "less good" student's side: you got it wrong—period. This is ideal for ranking students at scale.

By contrast, think about a subject like literature. It is arguably equally as important as mathematics, but much harder to rank, especially in its natural evaluation format: the writing and interpreting of essays, poems, stories, and so on. It has a much higher degree of *subjectivity*, and it is hard to distinguish submissions, at least in some strong sense. Of course, you can convert literature into something more objective—create multiple-choice questions, or ask someone to memorize and recite a poem from memory and grade for accuracy. But we all know this is not the heart of literature. These exercises do not capture its value.

For mathematics, by contrast, its objectivity seems for most people to be part of its identity. Its finality (its "rock-solidness," if you will) is taken to be its core feature. This is communicated in the culture of its standard pedagogy. For these reasons, it is possible to come to believe, as I think most people in society do, that mathematical ability is a clear and objective measure.

A chance encounter with a box of books

When I worked for the University of Chicago, I lived in Hyde Park, the neighborhood surrounding the university on the South Side of Chicago. Being a university neighborhood, it was not uncommon for a box of books to be placed on the side of the road with the word "Free" written on it. Drawn as I am to books, I would always stop and look.

On one particular day, I got especially lucky. I found a box of books that were part of "The University of Chicago School Mathematics Project." These were English translations of elementary-school math textbooks from all over the world. I happily gathered all of these books and brought them home.

The volumes that stood out to me were the Russian math texts. In particular, the Russian math texts had a type of question I had never seen before, and that seemed not to appear as frequently in

the other countries' books as I read. In addition to the usual "$4 \cdot 2 =$?" type exercises,[2] there were exercises in the other direction. I recollect reading a question something like:[3] Provide a question about a real-life scenario whose answer is $4 \cdot 2$." Or: "Make up a problem that uses the numbers 17, 10, and 5 and uses addition, subtraction, and multiplication."

I was blown away. I had no recollection of seeing anything like these "make up a problem" questions before in all my years of studying math (which I conservatively say was over 25 years at the time). I was so impressed by their possibilities that I started to incorporate questions of this structure into teaching my sons math, and even into my MBA classroom. Let me demonstrate why I find these so instructive with a simple example.

Four multiplied by two

When doing calculations in teaching my operations management classes, I started to take an aside to test how well my students really understand the concept of multiplication, using the "make up a problem" style questions that I learned from the Russian mathematics texts. My favorite one is:

Prompt: Provide a question about a real-life scenario whose answer is $4 \cdot 2$.

Take a moment before reading further to answer this for yourself. I provide a couple of lines for you to write your question:

Was it "easy" to provide your question? Let me share a few attempts I have heard over the years from various people, some quite confident in math. An answer I have heard a few times is:

Question 1: How many legs does a spider have?

Sound right to you? I find this one not super satisfying. When I think about a spider, the number 8 comes to mind for sure, but does $4 \cdot 2$? I think it is more $8 \cdot 1$. I count the legs one by one.

You may be thinking, "What's the difference between $8 \cdot 1$ and $4 \cdot 2$?" If so, that is an important lesson about multiplication for you. They are different even though they both *evaluate* to 8!

When I raise my objections to the spider-leg answer, the speaker often disputes:

"Well, I meant thinking about the spider in two sides?"

And then mimes in the air something that could be represented on paper as:

[2] In these Russian mathematics textbooks, multiplication was notated by a "$\cdot$" rather than a "$\times$".

[3] I sadly have since lost access to those texts in a moving debacle.

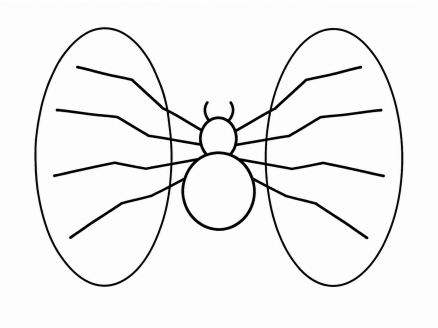

Figure 1 Counting the legs on a spider

But then I reply that this makes their original question not so perfect, which they typically concede. A "better" question is:

Question 1': How many legs does a spider have, thinking about it in two halves?

But I also do not like this question. It is a little awkward, is it not?

Another common question is something in the flavor of:

Question 2: I walk into the store and buy 4 candies, and my brother walks into the store and buys 4 candies. How many candies in total did me and my brother buy?

Do you like this one? I am still not super satisfied. A "better" answer to this questions seems to be $4+4$. Why? Because "I" and "my brother" are emphasized as being different. It makes me feel like it is a coincidence that each bought four candies, not that it was a pattern.

Why does this matter? Isn't $4+4=4\cdot 2$? Yes, they *evaluate* to the same thing, but they have different structural meanings. When I press on this point, the speaker often "improves" their question to something like:

Question 2': Two boys walk into a store. Each buys four candies. How many candies did they buy in total?

By "grouping" the two boys, it makes it look more like a pattern, making multiplication seem more appropriate.

I think maybe you are starting to see some of the complexities of the prompt: "Provide a question about a real-life scenario whose answer is $4\cdot 2$." Not only are there infinitely many possible "correct" solutions—I provided four above—but the quality of the answers is up for debate. Suppose we had to grade the solutions out of 10. Write down your grades in the space provided below:

Question 1 Grade: ___ (out of 10)

Question 1' Grade: ___ (out of 10)

Question 2 Grade: ___ (out of 10)

Question 2' Grade: ___ (out of 10)

I am guessing that you graded Question 1' higher than Question 1 and Question 2' higher than Question 2. But how about Question 1' versus Question 2'? Which did you grade higher? If the same, do you really believe both show equal understanding? What about the awkwardness of the wording of Question 1'? Should that matter? Who really thinks about spiders in "halves"? What does that even mean? Should this awkwardness be penalized? By how much?

As you can see, a lot of complexities and *subjectivities* arise. If you are someone comforted by the objectivity of math, this might be getting annoying for you. But can you say that this is *not* mathematics? In fact, I think most would have to admit that this is *deeper* mathematics than what most learned when doing multiplication as a child.

Reverse perspective

I want to take a short detour. I did not really have a name for this type of math exercise, the "come up with your own problem" exercise when I first discovered them.[4] But this style of question is amazingly versatile; I have used it to teach many things by now. As I thought more about its logic, though, I was searching for a name that might fit, and I came across a term from art.

When drawing something in three-dimensional space—basically, projecting three-dimensional space onto the plane—the most standard way is called *linear perspective*. In linear perspective, the lines converge inward towards a projection behind the picture. Think of how a road disappears into a point on the far-off horizon when depicted in the plane.

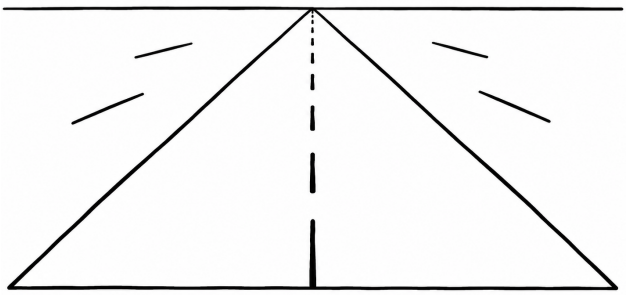

Figure 2 Linear perspective of a horizon

By contrast, *reverse perspective* has the objects depicted in a scene placed between the projection point and the viewing plane. I like to think about it as a fish-eye lens on a camera. Words are a little clumsy here, so let's look at a drawing in reverse perspective:

[4] I learned only recently that the area of "problem posing" in mathematical education covers some of this ground. See, for example, Brown and Walter (2005).

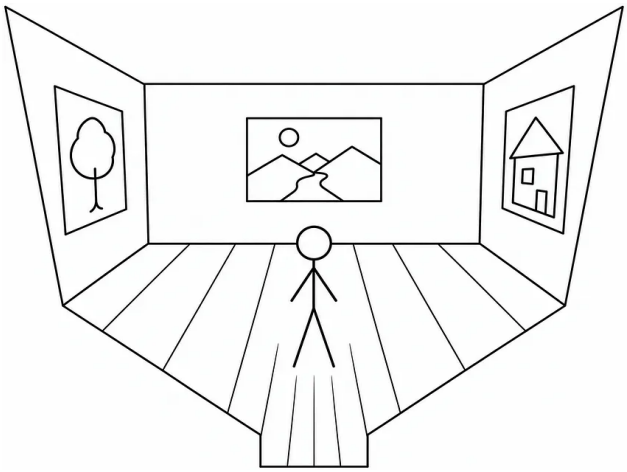

Figure 3 **The reverse perspective of a room scene**

To help set the contrast between the two perspectives even more, I borrowed a figure from Wikipedia:

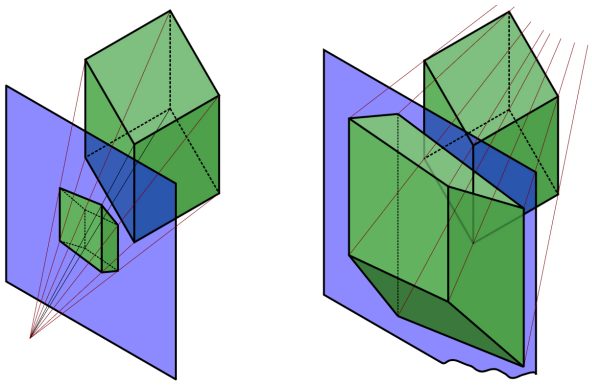

Figure 4 **Linear perspective (left) and reverse perspective (right). The blue plane is the "canvas", the green cube is the three-dimensional object to be captured in the "canvas", and the red lines denote the lines of projection. Figure: Rocchini Rocchini (2008), CC BY-SA 3.0.**

In linear perspective, the viewer looks *into* a represented world organized around a single viewpoint. In reverse perspective, by contrast, "the painting looks at the person, and not vice versa," and that "the world represented radiates towards those who open themselves to receive it" (Pirari 2024, p. 81).

The notion of reverse perspective, possibly coincidentally, also has connections to Russia. Art that employs *reverse perspective* was common in Byzantine and related medieval icon traditions,

including Russian icon paintings (Antonova 2010). In the early fifteenth-century Florentine Renaissance, however, linear perspective became the definitive system of Western pictorial space (Kemp 1992).

Let us return to the world of mathematics to see how this might be instructive.

For the most part, we have been taught mathematics in "linear perspective." We are drawn into mathematics. It has an objectivity and certainty. "Good" answers converge toward one logical truth. Math is not opened up to us, as it were; we are asked to come along with it. A typical mathematics exercise resembles linear perspective in this way: the student's reasoning is organized around an endpoint chosen in advance by someone else (a single answer, or a standard method that their work is expected to pass through).

A "divergent" mathematical task reverses this relationship. The mathematical idea becomes a kind of starting point from which many different questions, pictures, and ways of solving can spread outward, much as the lines of an icon spread outward toward the person looking at it. But we are not commonly asked to "create" mathematics, at least not until far into a mathematics education—which, for me, was essentially during my PhD. We are asked to submit to mathematics, to be drawn into its wholeness, its completeness.

This provides a geometric analogue for what psychologist J. P. Guilford distinguished as *convergent* and *divergent production*, now commonly discussed as convergent and divergent thinking (Guilford 1956). Convergent thinking directs possibilities toward a determinate solution, whereas divergent thinking generates a multiplicity of possible responses. Such divergence need not come at the cost of mathematical rigor. A "reversed perspective" exercise still asks for correctness (the spider does need to have eight legs), but it also leaves room for flexibility and originality.

Why does math emphasize convergence?

So you might ask: why are the "reverse perspective" questions, like those in the Russian math textbook, so uncommon? I hope a few of the reasons are obvious to you by now. Put simply, reverse-perspective exercises in math are hard to grade because they involve subjectivity. Did you enjoy grading Questions 1, 1', 2, and 2' earlier in this essay? Sound like fun to grade hundreds of these?

Linear-perspective math that yields a single verdict (right or wrong) is like track and field at the Olympics: there is no judgment involved in who wins. Reverse-perspective math is more like rhythmic gymnastics: we cannot escape its subjectivity. While there is undoubtedly, unquestionably, a "fastest man in the world," one can always debate who is the best rhythmic gymnast. In other

words, "linear perspective" mathematics is a "good" measuring stick, while "reverse perspective" mathematics is less so.

But it is worse than that. How do you teach someone to be good at reverse-perspective mathematics? It gets us closer to the question, "How do you teach someone to be a good poet?" The answer is not attending many lectures and drilling many problems. It is more like getting a lot of feedback, having a lot of interactions, and learning their own sense of "taste".

In OR/OM, there are some exercises that ask us to look at the world and express it in terms of math, not converging to a single answer, but using mathematics as a lens to see the world with greater richness. We call it "modeling" or "problem structuring". It is commonly thought to be more difficult, and more subjective, to teach than theorem statements and proofs. In comparison with "rigorous" math, modeling classes have a looseness that can bother many students, especially those entering the field from pure mathematics. It is harder to know when a model is "good", and more to the point, whether you are a "good" modeler.

Indeed, what makes a model "good"? It is hard to say, isn't it? Operations Management papers (in comparison to methodological papers more typical of Operations Research) have long introductions and model sections not because they are trying to pad the paper with fluffy content. Defending the value of subjectivity in a model often takes much more verbosity than a "QED".

I know there will be readers of this essay who do not like the idea of subjectivity in mathematics. I know it would have bothered my high school math teacher, who felt math's objectivity was one of its redeeming features. Luckily for these readers, we can be well assured that math will never be as subjective as poetry.

Mathematics with a proof and problem-solving calculator

There are many news articles, social media posts, and bouts of soul-searching among mathematicians and applied mathematicians these days. AI has become a powerful tool in proving theorems and solving mathematical problems. I can attest to this personally, as well as hearing almost daily recently from mathematical friends amazed at the power of modern AI tools. What does it mean that mathematics now has AI tools that are better than almost all real people at linear-perspective mathematics? Give AI a question with "an" answer, T or F, and it has a very good—and rapidly accelerating—chance of providing the answer.

Why is AI so good at linear-perspective mathematics? The goal of this kind of mathematics is very objective and verifiable: prove true or show false. And the "move set" for deciding true or

false is also quite small. Logic needs to follow clearly established steps that are even now being formulated into something like a computer language, "Lean". Linear-perspective math looks a lot more like chess than poetry.

Anyone with a child knows how hard it is to get students to multiply big numbers by hand instead of using a calculator, unless they turn it into some kind of game or invoke the fear of appearing to be stupid. But now we have a 'calculator' that can solve an astonishing range of word problems, contest questions, and other convergent mathematical tasks that occupy most students' mathematical education. How can we motivate them to learn math with such powerful tools at their fingertips? An old question with new twists. We might have told kids we need to learn to do arithmetic by hand to be able to make sense of harder problems that come later. But when do these "harder" problems come now with access to an AI "calculator"? Will they ever come for any but the most immersed in mathematics?

Luckily, we know that this is not all of mathematics. There is a lot of openness and creativity in mathematics—that almost nobody in our society has ever experienced.

Descriptive mathematics

Like most people mentioned above, my skills in converging toward the right answer have been surpassed by AI. I will admit the following: I was surpassed already many moons ago. Probably sometime in 2024, I could see that ChatGPT was better than me at attempting to prove math results and solve math problems. I saw the writing on the wall then, in part because I wanted to see it.

I was ready to be free of mathematics my whole life, and it finally looked like I would be free forever. Except … I needed to teach math to my sons.

I thought about what I would teach them. I could think of three types of mathematics, and I explained all three types to them. The first is *school math*: what you learn in elementary and high schools in Canada, where I learned it. Its primary lessons are precision and obedience. It is not a place to ask questions, unless we as a society somehow change it. It is linear-perspective mathematics through and through.

The second is *competition math*, which asks for more creativity, but itself has its own "hidden" tropes and techniques that you could learn outside of school (but typically not in it) for a high price. But, possibly even worse than school math, it only cares about the answer, not the process. Most math competitions I have seen (and most people see) are multiple-choice or graded on the final answers alone. For the most part, it is linear-perspective mathematics *par excellence*.

But I identified a third type of mathematics: *descriptive mathematics*, which I have been personally teaching my sons and some of their friends. Descriptive mathematics is reverse-perspective mathematics, native to AI as a powerful calculator. The goal is to use the tools of mathematics as a language to describe the world.

I have taught them, so far, set theory, graph theory, basic optimization theory, and now we are moving on to some probability. I do not prove things or calculate things; I let AI do those. But we are not happy with answers from AI. The answer is not the point. We need to put those ideas into our words and into our world, the world we see, the wide world in reverse perspective. The program is full of exercises like: "Provide a question about a real-life scenario whose answer is $4 \cdot 2$."

* * *

* * *

Where do we go from here?

The future of mathematics is bright, but for it to be bright, we need to be braver about it being open: a future of more subjective mathematics. I believe this type of education, going back to the earliest days of mathematics education, would bring more people to the path of mathematical modeling of the world, and ultimately to many OR/MS/OM classrooms.

In that world, mathematics will be a less powerful tool for ranking students, and it will be a subject much harder to grade. But this is a path to making the subject more human. If mathematics is robotic, we know it will be automated. But we know that the true spirit of mathematics is indeed human, so I am not worried. Let's take a reverse perspective and see many possibilities opening up to us. The subject of mathematics has a lot of room in it: room for more people, maybe even room for everyone.

References

Ahuja RK, Orlin JB (2001) Inverse optimization. *Operations Research* 49(5):771–783.

Antonova C (2010) On the problem of "reverse perspective": Definitions east and west. *Leonardo* 43(5):464–469, URL `http://dx.doi.org/10.1162/LEON_a_00039`.

Brown SI, Walter MI (2005) *The Art of Problem Posing* (Mahwah, NJ: Lawrence Erlbaum Associates), 3 edition.

Daw A, Gupta V, Rusmevichientong P (2026) Modeling first: Rethinking undergraduate operations management with AI. *INFORMS Transactions on Education* URL `http://dx.doi.org/10.1287/ited.2025.0168`, articles in Advance, published online 11 May 2026.

Grossman TA Jr (2001) Causes of the decline of the business school management science course. *INFORMS Transactions on Education* 1(2):51–61, URL `http://dx.doi.org/10.1287/ited.1.2.51`.

Guilford JP (1956) The structure of intellect. *Psychological Bulletin* 53(4):267–293.

Kemp M (1992) *The Science of Art: Optical Themes in Western Art from Brunelleschi to Seurat* (New Haven: Yale University Press).

Little JDC (1961) A proof for the queuing formula: $L = \lambda W$. *Operations Research* 9(3):383–387.

Morris WT (1967) On the art of modeling. *Management Science* 13(12):B707–B717.

Pirari G (2024) Perspective: Realism or illusion? *Philosophical Letters. Russian and European Dialogue* 7(2):71–87.

Powell SG (1995) The teachers' forum: Teaching the art of modeling to MBA students. *Interfaces* 25(3):88–94, URL `http://dx.doi.org/10.1287/inte.25.3.88`.

Powell SG (1998) The studio approach to teaching the craft of modeling. *Annals of Operations Research* 82:29–47, URL `http://dx.doi.org/10.1023/A:1018918921334`.

Rocchini (2008) Reverse perspective. Wikimedia Commons, URL `https://commons.wikimedia.org/wiki/File:Reverse_perspective.svg`, sVG image, licensed under CC BY-SA 3.0.

Zhang P (2026) Teaching applied optimization with large language models. *INFORMS Transactions on Education* URL `http://dx.doi.org/10.1287/ited.2024.0096`, articles in Advance, published online 12 May 2026.

Appendix. Some "reverse perspective" exercises in OR/MS/OM

The questions below are grouped loosely by the courses in which they might appear: optimization, probability, and operations management. The operations management questions draw on themes from *Operations, Out Loud*, the course book on operations management I am writing with Nanxi Zhang; the others I have drafted for this appendix.

Each question is followed by a "linear-perspective cousin"—a question of the kind found in most textbooks that exercises the same idea—and then by some discussion of what changes when the direction of the question is reversed. In the discussions I have tried to show a few answers of the sort students give, and to say what I find more and less satisfying about them, in the same spirit as the spiders and candies of the essay. I expect readers will grade these answers differently than I do, and that is, in a sense, the point.

One more remark before beginning. Any of these prompts can be typed into an AI tool, and it will return a perfectly respectable answer in a few seconds. I do not think this defeats the exercise. As with the $4 \cdot 2$ question, I believe the learning happens less in producing a first answer and more in the conversation that follows, about whether the answer is any "good" and why. A student who asks an AI for three answers, and can then explain which is best and what is lacking in the other two, has (to my mind) done the exercise.

A1. A linear program whose answer is $(4, 2)$

Prompt: Describe a real-life decision that can be written as a linear program in two decision variables, x and y, whose optimal solution is $x = 4$ and $y = 2$. Say what x and y are, and what each constraint means.

Linear-perspective cousin: "Solve the following linear program graphically: maximize $3x + 4y$ subject to $x + y \le 6$, $x + 2y \le 8$, and $x, y \ge 0$."

The cousin and the prompt are, in a sense, the same problem run in opposite directions. In the cousin, the constraints and the objective are given, and the student converges on the corner $(4, 2)$. In the prompt, a solution is given, and the student must build a small world around it.

The first answer I tend to hear is the linear-programming version of the spider:

Answer A: A shop can sell at most 4 chairs and at most 2 tables a day, and makes \$1 on each. Maximize $x + y$ subject to $x \le 4$ and $y \le 2$.

It is correct, but like the spider's legs it feels to me like $8 \cdot 1$ dressed up as $4 \cdot 2$. Each variable is fenced in by its own constraint and the two decisions never interact, so there is really nothing to optimize (nobody needs a linear program to decide to make as many of everything as they are allowed). A "better" answer makes the decisions compete for something:

Answer B: A small bakery makes trays of muffins (x) and trays of scones (y). It has 6 hours of oven time, and each tray of either kind takes an hour. It has 8 kilograms of flour; a tray of muffins uses 1 kilogram and a tray of scones uses 2. A tray of muffins earns \$3 and a tray of scones \$4. Maximize $3x + 4y$ subject to $x + y \le 6$ and $x + 2y \le 8$.

Here the oven and the flour are both used up at $(4, 2)$, and the answer comes out of the tension between them. To write Answer B, a student must understand, at least intuitively, that the optimal corner sits where two constraints meet and that the objective has to "point" in the right direction between them; if scones earned \$10 a tray instead of \$4, the bakery would make only scones. These are ideas the cousin tests only indirectly, if at all. A student can solve the cousin by evaluating every corner, as a calculator would, without ever seeing why the answer sits where it does.

How should we grade Answer A against Answer B? Both are correct. Is a story with three constraints better than one with two? Is a story about a bakery better than one about a refinery, because the student has actually been inside a bakery? That is an interesting consideration I don't want to litigate here.

Readers might also recognize the prompt as a classroom relative of *inverse optimization*, which asks what objective would make an observed decision optimal (Ahuja and Orlin 2001).

A2. An unbounded linear program

Prompt: A group of analysts modeled a planning problem using linear programming. When they implemented their algorithm, the problem ended up being unbounded. Create a real-life setting where this might have reasonably arrived.

Linear-perspective cousin: "Determine whether the following linear program is unbounded: maximize $2x + y$ subject to $x - y \leq 1$ and $x, y \geq 0$."

In most textbooks, unboundedness is presented as a property of a system of inequalities, a kind of pathology to be diagnosed. The cousin can be answered by finding a direction in which one can walk forever while the objective keeps improving, and a student can do this without ever asking what x and y are.

The prompt asks for something rather different, because nothing in the world is unbounded (no business earns infinite profit). An unbounded linear program is therefore always a statement about the model rather than about the world. To answer the prompt, the student has to invent a model that is wrong in a particular way and then say how. A typical answer goes something like this: a trader buys a product in one market and sells it in another at a higher price, and the model lets her buy as much as she likes. What did the modeler forget? Money to buy with, perhaps, or space to store the goods, or the fact that prices move when someone buys a lot of something.

What I like about this prompt is that the mathematics and the world get to check each other. An answer of "infinity" is, in effect, a message from the model about itself, and perhaps the most interesting part of a student's response is the (potentially ambiguous) list of things that were forgotten.

A3. An infeasible linear program

Prompt: A group of analysts modeled a planning problem using linear programming. When they implemented their algorithm, the result ended up infeasible. Create a real-life setting where this might have reasonably arrived.

Linear-perspective cousin: "Determine whether the following linear program is feasible: maximize $x + y$ subject to $x + 2y \leq 4$, $3x + y \geq 15$, and $x, y \geq 0$."

In most textbooks, infeasibility is presented as a set of constraints that cannot all be satisfied at the same time. The cousin can be answered by drawing the constraints and seeing that the regions they allow do not overlap, or by letting the software report that no solution exists. A student can do this without ever asking what x and y are, or who wrote the constraints, or why.

The prompt asks the student to imagine the people behind the model, and here infeasibility differs from unboundedness in an interesting way. An unbounded model is essentially always wrong about the world, since nothing in the world is unbounded. An infeasible model, by contrast, may be right about a tough situation, or it may simply be wrong. Some answers find a mistake in the model: a constraint typed with the inequality pointing the wrong way, one constraint measured in hours and another in minutes, or a rule written as if it could never be broken when, in practice, people

bend it all the time. A nurse schedule that must give every nurse every day off they asked for, and also cover every shift, may simply have no solution. Other answers find that the model is right and the world is asking for too much: a sales team has promised customers more units this month than the factory can make, and the model is the first place where anyone has had to add up all the promises at once.

The more interesting part, for me, comes after the setting is chosen. Once the student has a setting, the natural question is what the analysts should do about it. Some constraint has to give, and the mathematics alone cannot say which one. Should the factory miss the delivery date, or run overtime, or disappoint a smaller customer to keep a larger one? Deciding which constraint to loosen is a decision about who bears the cost, and that is a conversation the analysts will have to take back to the people who gave them the constraints in the first place. A student who finds the typo and a student who finds the over-promising sales team have both answered correctly. The second student, I think, has found the more interesting problem, though I would have a hard time explaining to the first why that should earn more marks.

A4. A shortest path that is not the fewest steps

Prompt: Describe a real network—of roads, people, pipes, or anything else—in which the shortest path from one point to another is not the one that takes the fewest edges. Say what "short" means in your network.

Linear-perspective cousin: "Use Dijkstra's algorithm to find the shortest path from node A to node F in the network below."

The cousin hands the student the lengths of the arcs, whereas the prompt asks the student to decide what "length" should mean in the first place. The most common answer is a drive across a city, in which the highway route has more segments but takes less time than the direct route through downtown. This is a fine answer, but more interesting answers come when "short" is not about distance or time at all. Consider someone hoping for an introduction to a hiring manager, who may prefer a path through three people who know one another well over a path through a single acquaintance who barely remembers them. Here the length of a path is something like the risk that the message dies along the way.

At this point the student has run into a question that the cousin never asks: do the lengths add up along a path? Travel times add up. The chance that a message survives three handoffs, however, is the *product* of three probabilities rather than their sum, so it is not obvious that the shortest-path model fits at all. A student who knows a little more may notice that taking logarithms turns products into sums, and that the model can be rescued. I would struggle to rank the commuter's answer against the job-seeker's. The commuter's is cleaner, and I find the job-seeker's richer (if a little more fragile).

A5. A story for a picture

Prompt: Below is a network. Describe a real-life setting that this network could represent. Say what the nodes are, what the arrows are, and what the numbers mean, and make sure your story explains every part of the picture.

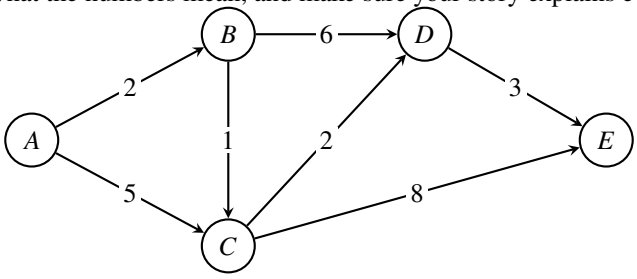

Linear-perspective cousin: "In the network above, the numbers are travel times in minutes. Find the shortest path from A to E."

This prompt turns A4 around once more. In A4, the student was given an idea, a short path, and asked to find a network in which it shows up. Here the student is given the network itself, with every node, arrow, and number already fixed, and is asked to find a world it could describe. The cousin uses the same picture but does not investigate how it maps to the real world; it refers only to itself. Once the numbers are called minutes, there is nothing left to decide except the route. The quickest route goes from A to B to C to D to E and takes 8 minutes, even though it uses four arrows, while the route through C alone uses only two arrows and takes 13 minutes.

The prompt, by contrast, leaves all of this open. Students usually begin with roads, because roads are what networks look like to most of us. Then the arrows start to cause trouble. Why can a car go from B to C but not from C back to B? One-way streets will do, but the student has to say so. Other students leave roads behind and choose something where direction comes more naturally: water running downhill through pipes, courses where one must be taken before another, the steps of a recipe, or the chain of people a rumor passes through. Each choice changes what the numbers can mean. For pipes, a number is most naturally how much water an arrow can carry. The choice of units for the numbers is also needed to tell the story effectively.

Some stories explain most of the picture but not all of it. A student might describe a delivery route but have nothing to say about why there is no arrow from D back to B, or why the number from C to E is so large. Should a story that accounts for every arrow be graded above a more vivid story that leaves one arrow unexplained?

A6. One half times one third

Prompt: Provide a question about a real-life scenario where the probability of an event happening would be $\frac{1}{2} \cdot \frac{1}{3}$.
Linear-perspective cousin: "A fair coin is flipped and a fair six-sided die is rolled. What is the probability of getting heads and a number less than 3?"

This prompt is deliberately the closest to the essay. Many students answer it by turning the cousin around, and their answer is correct. But the coin and the die are, I think, the spiders of probability: gadgets built so that the multiplication rule holds. Rolling a single die and asking for the chance of a 1 gives $\frac{1}{6}$ as well, and is the $8 \cdot 1$ of this problem: the "right" value, but with the wrong structure.

When students reach for real life instead, the conversation changes. "The probability that it rains tomorrow is one half and the probability that my bus is late is one third, so the probability of both is one sixth." Is it? Rain makes buses late. The multiplication now rests on an assumption, independence, that the student must defend, and in real life that defense is rarely airtight. Another student might offer the first customer at a café, who is a woman with probability about one half and orders tea, one of three drinks on the menu, with probability one third. Here there are two hidden assumptions rather than one: that the two events are independent, and that the three drinks are equally popular.

A7. Stocking above, and below, average demand

Prompt: Describe two real businesses that sell a perishable product with uncertain demand: one in which the right amount to stock is well above average demand, and one in which it is well below. Explain what makes the two different.

Linear-perspective cousin: "Daily demand for newspapers is normally distributed with mean 100 and standard deviation 20. Each paper costs $1 and sells for $3, and unsold papers are worthless. How many papers should be stocked?"

The cousin can be done by placing numbers into the critical-ratio formula; the student needs to know where the numbers go, but not why. The prompt gives away the conclusion—stock above the mean, stock below the mean—and asks for the world. To answer it, a student must have understood that the decision turns on comparing the cost of having one too few with the cost of having one too many, and must find real businesses in which that comparison tips each way.

What makes this prompt subjective is that the costs in real businesses are not printed on a card. A student might argue that a pharmacy should stock well above expected demand for flu shots, because a customer turned away may never come back, while a grocery store's sushi counter should stock below expected demand, because unsold sushi is thrown out every night. Another student might argue that the lost flu-shot customer simply goes next door and is no great loss. Who is right depends on things neither student can easily measure. Both have understood the newsvendor model better than a student who computed the cousin's answer perfectly.

A8. Sleuthing with Little's Law

> Prompt: Describe a real-life situation in which Little's Law lets you estimate something you could not easily observe directly. Say what is flowing, where the boundary of your system is, and which two quantities you would measure.

Linear-perspective cousin: "Patients arrive at a clinic at a rate of 10 per hour, and each spends an average of 30 minutes in the clinic. On average, how many patients are in the clinic?"

The cousin hands over two of the three quantities in Little's Law (Little 1961) and asks the student to multiply. Almost every student can do this, and many can do it without any sense of why the law should be true or when it would be useful.

The prompt asks the student to find the use. The first answers often recreate the cousin—a restaurant, a clinic—where all three quantities are easy to see and the law does no real work. Better answers find a system in which one quantity is hidden. A university program has about 400 students enrolled at any time and admits about 100 a year, so students take about 4 years on average to pass through, even though nobody followed any of them from entry to exit. To write such an answer, a student has to choose a boundary (does a student on leave count as "inside"?), choose what is flowing, and decide whether the system is settled enough for the law to apply. These are all judgment calls.