UCLA
UCLA Previously Published Works

Title

The Operations of You: How Bottlenecks, Queues, and The Toyota Way Explain Why Time Management Advice Works

Permalink

https://escholarship.org/uc/item/2m41j7kv

Journal

Management and business review

ISSN

2694-1058

Authors

Corbett, Charles
Roels, Guillaume
Yoo, Onesun Steve

Publication Date

2026-11-08

Peer reviewed

eScholarship.org Powered by the California Digital Library
University of California

The Operations of You: How Bottlenecks, Queues, and The Toyota Way Explain Why Time Management Advice Works

Forthcoming in *Management and Business Review*

Charles J. Corbett, UCLA Anderson School of Management
Guillaume Roels, INSEAD
Onesun Steve Yoo, UCL School of Management, University College London

November 8, 2026

Executive summary

Readers of *Management and Business Review* will be familiar with much of the advice encountered in the time management literature. “Don’t focus on the urgent at the expense of the important.” “No is a two-letter word.” “Keep your workspace organized.” In our experience with executives, colleagues, MBA students, and all high-functioning professionals, the adoption of these practices is quite limited. We attribute this to a lack of understanding of their complementary nature. Each recommendation is typically considered in isolation, but in fact they all amplify each other. To make this case, we draw on three central pillars of operations management (OM), and offer new perspectives on time management in doing so. This perspective also clarifies how each pillar paves the way for the others. The consequence of this complementarity is that, over time, the value of many of these practices that we have all heard about is even greater than might appear from the way they tend to be portrayed. We hope that thinking about “the operations of you” will provide additional stimulus to try it out, one small change at a time.

Introductory Note

The authors show how applying insights from operations management to personal time management can help explain why many common time management practices work, and how they are complementary. They argue that our time is the bottleneck, that some slack is necessary, and that the cumulative effect of implementing small continuous improvement steps can be powerful and self-propelling.

Readers of *Management and Business Review* will be familiar with much of the advice encountered in the time management literature. “Don’t focus on the urgent at the expense of the important.” “No is a two-letter word.” “Keep your workspace organized.” The list goes on. Much of the advice is sensible. Nevertheless, in our experience with executives, colleagues, MBA students, and all high-functioning professionals, the adoption of these practices is quite limited. We attribute this to a lack of understanding of their complementary nature. Each recommendation, while plausible, is considered in isolation, while, in fact, these practices all amplify each other. Any one of them, by itself, may not seem worthwhile, but together they provide a powerful self-perpetuating force. To make this case, we draw on three central pillars of operations management (OM). That might seem a strange place to look for advice on human time management, but OM revolves around getting the most out of scarce resources. If our time is scarce, the core principles of OM should apply to us too. We translate each of these pillars to the realm of personal time management. By drawing on tools and insights from OM, this translation offers new perspectives on time management. It also clarifies how each pillar paves the way for the others. The consequence of this complementarity is that, over time, the value of many of these practices that we have all heard about is even greater than might appear from the way they tend to be portrayed.

The first pillar of OM is the Theory of Constraints as popularized in Goldratt’s widely-read novel *The Goal*;[1] its central prescription is to identify the bottleneck in an organization, and then manage accordingly. If our time is our main constraint, we ourselves are the bottleneck in our own lives. (Don’t worry, that is a good thing: the bottleneck should be the most valuable resource.) The time-tested lessons from *The Goal* should apply to us too. Do you have enough time to do everything you want? If you do, you can stop reading this. If not, we hope that taking 5-10 minutes to read this short article will be an excellent investment of your time.

If our time is our main constraint, we ourselves are the bottleneck in our own lives.

The unpredictability of what lands on our plate, and when, poses an inescapable challenge, and brings us to the second pillar of OM: queueing theory. Most readers have a longer to-do list than they like, and it takes too long to get things done. Queueing theory highlights the combined effects of utilization (being busy) and variability. It explains why some slack is essential to tame the to-do list, and why reducing variability, which may seem only a second-order effect, is so valuable. It also teaches that even small improvements matter. So how does one make those small improvements?

Enter the third pillar: the discipline of continuous improvement, as embodied in Liker's book *The Toyota* Way.[2] The Toyota Production System, the operational embodiment of the philosophy that defines the Toyota Way, is the gold standard, proven over decades around the world, at Toyota and copied by many others, in both manufacturing and service industries. The improvement from any single change is often modest; the cumulative effect of relentlessly pursuing continuous improvement makes it so powerful. Many practices associated with the Toyota Way apply to our own time management too, as does the power of compounding: imagine trying to match the productivity of a colleague who has made one incremental improvement every week for the past 30 years.

In this paper, we explain how these three pillars of OM apply to ourselves. The specific time management practices we refer to have all been advocated elsewhere, but often in a rather ad hoc manner. Revisiting the abundant time management advice through this OM perspective demonstrates how many individual recommendations reinforce each other. One student, a former software engineer, realized "time management is just operations."[3] We hope that thinking about "the operations of you" will provide additional stimulus to try it out, one small change at a time.

You are the bottleneck; now what?

Goldratt's *The Goal* introduces a five-step process for managing bottlenecks. Let's look at how applying these to ourselves helps us make the best use of our time.

Step 1: identify the bottleneck

A bottleneck is any resource whose capacity is less than or equal to the demand placed upon it. Consider a small clinic, with one physician who can see ten patients per hour, and a receptionist who can check in fifteen patients per hour. If the receptionist actually admitted fifteen patients per hour the waiting room would fill up: the physician is the bottleneck. A bottleneck is not a flaw, it is an inevitable feature of any constrained system. Ideally, the bottleneck should be the most scarce resource. If the clinic's only constraint were chairs for patients to sit on, you could easily alleviate that bottleneck by buying more chairs.

What is the bottleneck in your life? Your time. Cognitive neuroscientist Levitin implicitly treats the mind's central executive as the bottleneck resource, emphasizing its capacity constraint. [4] His exploration of how to best use that scarce resource provides a deeper scientific underpinning for several of the parallels between OM and time management.

Step 2: optimize the bottleneck

Goldratt's overarching prescription is to avoid wasting bottleneck time: "an hour lost on a bottleneck is an hour lost for the entire plant." Applied to us, that can take many forms.

- **Routines.** These allow you to do more, without adding raw hours of work time. Former Apple CEO Tim Cook tried to always eat at the Apple cafeteria so he did not have to think about what to eat.[5] An attorney observed that "by creating better systems and routines, I can manage my workload more effectively and keep things moving without becoming overwhelmed."

- **Have a to-do list ready.** A bottleneck should not spend time not knowing what to do. You do not have to do everything on the list immediately, but whenever an idle moment presents itself the list will give you useful things to prioritize.
- **Avoid harmful multitasking.** Every time we switch between activities, we incur an invisible but real cognitive switching cost. Asking the brain to switch attention between activities burns up fuel, leading to mental exhaustion.[4] Our mind cannot actually multi-task; it rapidly jumps back and forth between tasks, incurring switching costs every time. Multitasking is particularly bad when switching between two cognitively demanding tasks.[6] Newport emphasizes the importance of working in larger batches of time,[7] just as Goldratt already recommended for bottlenecks. One simple but powerful tool is to switch off (almost) all notifications on your devices. 'Focus' and 'Do Not Disturb' modes can silence notifications automatically during set hours, blocks of deep work, or at specific locations. Even if you don't respond to every notification, your mind wastes scarce capacity observing and then trying to ignore it. Of course, humans have to work on multiple tasks during the day; the key is to always do so in such a way that switching is minimized and that the cognitive cost of the remaining unavoidable switching is kept as small as possible.
- **Outsource.** Even if you are good at something, that doesn't mean you should be doing it; and if you're not good at it, then even more reason to outsource to someone who is. AI has vastly expanded the scope and feasibility of what we can outsource. Sometimes you can use money to "buy time" by outsourcing work; happiness and time experts such as Holmes[8] and Whillans[9] advocate this wherever possible.
- **Say no.** "No is a two-letter word" is one of the most frequently-encountered bits of time management advice. Is this task important enough to do? What will happen if we just don't do it? Drucker distinguishes between doing the right things (being effective) and doing things right (being efficient).[10] Saying no is about effectiveness; we examine efficiency later. As Porter and Nohria point out, CEOs cannot create more time, so they have to be careful how they allocate it.[11] Saying no can also mean deleting apps from

your phone that take more time than they are worth, minimizing the amount of time spent on commuting, and much else.

A key element is recognizing the opportunity cost of your time. Saying "yes" to something inevitably means saying "no" to something else. A common fallacy is to embark on one project that sounds more appealing than another, but failing to recognize that it will take much more time, so the value per unit bottleneck time may be lower. A busy independent consultant may well be better off taking on a new engagement yielding $50,000 for 50 hours of work than a seemingly more juicy $100,000 project that will take 250 hours. Decisions about which tasks to pursue should be made based not only on the tasks, but also on the opportunity cost of your time. To make matters worse, that opportunity cost is often increasing over time, as the consultant gains more experience or an executive's career advances. That opportunity cost need not be purely financial; for a young professional, starting a family may also drastically increase that opportunity cost. All this makes investing in time management today all the more urgent, as we explore later.

Once we have optimized our bottleneck, what next?

Step 3: subordinate everything else to the bottleneck

The other part of Goldratt's prescription about making optimal use of the bottleneck is to stop worrying about non-bottleneck time.

- **Delegate.** Is this task something only you can do? A common fallacy is "I'll do it myself because I can do it much faster than my assistant." This ignores the difference in opportunity costs: if yours is $1,000 per hour, and your assistant's is $100, it is worth delegating even if your assistant takes five times as long. If you worry that it will take you longer to explain the task than to just do it yourself, it may be time to write out the steps involved so you can delegate more efficiently.

- **Have all inputs ready when you need them.** If you need to make a decision that depends on others, make sure you have their input. If you start weighing pros and cons and some pieces are missing, you've just wasted bottleneck time.
- **Non-bottlenecks should not keep you waiting.** When you take your (non-bottleneck) teenager to soccer practice, they should be waiting for you, not the other way around. Idle time on a non-bottleneck is perfectly fine (and actually has other benefits too).
- **Avoid pseudo-productivity.** Reassure people working for you that it's OK to not always try to look busy (what Newport calls pseudo-productivity[12]), otherwise they will generate work for everyone, including you.

Individuals working in organizations need to recognize the limitations to this mindset. If everyone follows our prescriptions rigidly, treating their own time as the ultimate bottleneck, organizations will grind to a halt. It will become impossible to find meeting times; teams will stop functioning. Organizations need to establish a hierarchy of bottlenecks; the CEO is usually the dominant bottleneck, so the rest should be subordinated to the CEO.[11] The relative ranking of bottlenecks need not correspond perfectly to the organizational hierarchy; a single specialized engineer can be a major bottleneck, as illustrated nicely in *The Phoenix Project*,[13] which applies the theory of bottlenecks to the world of IT. Managers need to identify the appropriate organizational rhythms, balancing between respecting everybody's individual time, while enforcing the degree of synchronization that the organization needs.

After you have optimized your time and others are not interfering with it, what next?

Step 4: elevate the bottleneck

You cannot squeeze 25 hours in a 24-hour day, but you can increase efficiency (more output for given capacity) and cognitive capacity itself. Increasing efficiency is the "get more done" mindset that is so common in time management.

Although you cannot create time, you can create more useful time. Preventive maintenance is well-established in OM: sometimes it is better to shut down a bottleneck machine to help keep it running optimally, rather than waiting for it to fail and face a lengthy downtime to fix it. Practices such as exercise, diet and mindfulness are all forms of preventive maintenance, that help prevent burnout and increase how much cognitive capacity you have during the day. According to a former military executive, "thinking about personal productivity through an operations lens makes mindfulness feel much more practical and actionable." In fact, such practices may allow you to experience more time than what your synchronized mobile devices might suggest.[14] Elevating the bottleneck may sound like a heavy lift, but does not require big steps; many small increments add up.

Although you cannot create time,

you can create more useful time.

Step 5: re-optimize

Managing by the bottleneck is a never-ending cycle. Traditionally, Step 5 is "find the next bottleneck". Once you have elevated the bottleneck, the next weakest link becomes the new bottleneck. If the clinic we saw earlier hires an additional physician, the receptionist becomes the bottleneck. Now we have to optimize the receptionist's time, subordinate everything to the receptionist, elevate the receptionist's capacity, etc.

Our bottleneck is not going to shift from our time to something else, but we still need to keep iterating through these steps, because circumstances change. New technologies come along that can help or hinder; we get assigned to new projects and have new co-workers; our priorities and objectives may shift. This requires us to go back to Steps 2, 3, and 4, and re-optimize as circumstances change, rethink how we coordinate, and explore new ways to elevate our capacity.

Some recent time management literature, such as Burkeman's *Four Thousand Weeks*,[15] focuses more on wellbeing rather than productivity for its own sake, and we couldn't agree more. Even Allen's *Getting Things* Done,[16] one of the all-time best-selling books on time management, is subtitled *The Art of Stress-Free Productivity* and emphasizes value and wellbeing. The ultimate objective of managing time is not to just get more done, but to get more value and joy out of what you do. A financial analyst at an entertainment firm found that "thinking about mindfulness this way makes it feel less like a separate habit and more like good operational discipline for how we manage our time and attention." That requires paying attention to what you care about, and why, but that takes time. It requires taking care of physical and emotional health, but that also takes time. As we constantly learn more about our true values, we inevitably need to re-optimize. Managing by the bottleneck is a cycle, not a checklist.

The first major pillar of OM, the Theory of Constraints, essentially tells us how to get the most out of our time. What happens when bottlenecks collide with uncertainty?

The need for slack

The traditional bottleneck is a machine, and the goal is to use it optimally, often thought of as striving for 100% utilization, but that ignores variability. Think of yourself as a "server." Tasks arrive at random times, each of which will demand an uncertain amount of your time (your capacity) to complete. This is a basic queueing system. Queueing theory describes how long your queue, or your to-do list, will be, and how long it will take you to complete a task after it joins your queue. Even though your personal life is far more complex than a simple queueing model, the central insight from queueing still holds: the to-do list and the completion times increase disproportionately with utilization and variability. Utilization is how busy you are. Our physician could see ten patients per hour. If only one patient actually shows up per hour, the physician's utilization is 1/10 = 10%, and there will almost never be a queue. With so much slack, variability in when patients arrive and how much time they need doesn't really matter. If

eight patients arrive per hour, on average, the physician's utilization is 80%. Even though their capacity is sufficient on average, there will often be queues due to the variability in patient arrivals and service durations. If one patient takes 20 minutes, several more patients might arrive during that time. Eventually, though, the queue will clear up.

If ten patients arrive per hour, the queues will grow infinitely long. The physician is the bottleneck. Sometimes they have nothing to do, due to the variability in arrivals. The physician only has precisely enough capacity to meet demand, so whenever they lose some time, they can never make it up. A former software engineer noted that "even if I planned my day perfectly around average expected timing, any deviation could result in the entire schedule being messed up." Queueing theory shows that even with utilization below 100%, variability creates queues and delays. The closer you get to 100% utilization, the worse it gets. This is the "curse of variability" in Figure 1: getting just a little busier can cause queues and delays to become disproportionately longer. The effect of small changes is a fundamental insight from queueing theory. If you are already very busy, a few unexpected tasks can completely throw your fragile schedule into disarray, causing massive delays and all the associated headaches. In Figure 1, when your utilization increases from 85% to 90%, the effect on your to-do list is much more dramatic than if it increases from 50% to 55%, and the larger the variability, the higher this difference. An analyst at a healthcare company recognized this: "When my schedule is too full, even small disruptions create a backlog that feels much bigger than it should."

On the positive side, if you're already very busy, then even a relatively small decrease in utilization, e.g. from 95% to 90%, can lead to noticeable improvements in your inbox. How do we make that happen? Because utilization and variability conspire to cause these delays, those are the two levers we can use to make things better.

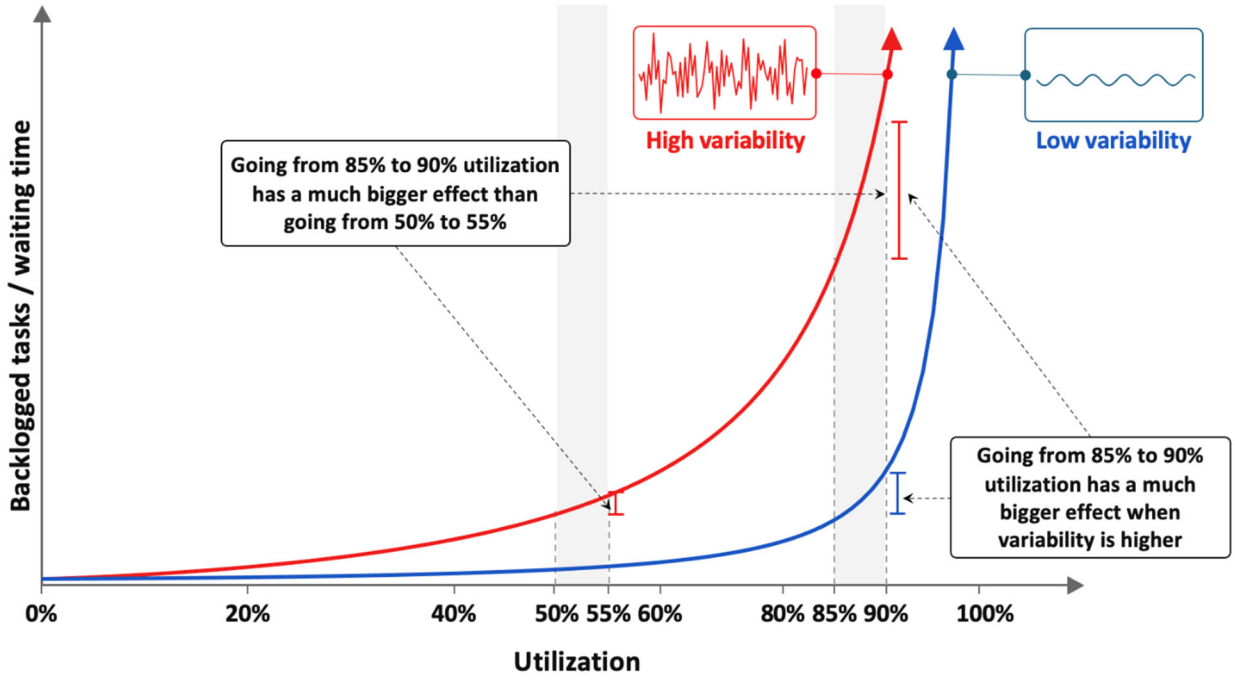


Figure 1: A backlog grows disproportionately with utilization and variability

Lever 1: reduce utilization (create slack)

This is what we discussed earlier, under optimizing and elevating the bottleneck. By reducing arrival of work (saying no), or by working faster (e.g., using routines), you decrease your utilization. This will immediately shorten your to-do list and reduce delays, which has further positive spillovers. It is stressful to carry around a long list of to-do's, so your cognitive capacity decreases, so it takes you longer to complete tasks, so your to-do list becomes even longer, etc. The vicious cycle is clear, as is the remedy: limit how many items can be on your to-do list, such as by finishing tasks before accepting new ones. We recognize that this is an oversimplification: saying no will not automatically free up time on your calendar, as something else will take its place, if you let it. Truly creating some slack requires taking to heart the inescapable truth in the aforementioned Burkeman's *Four Thousand Weeks*: there will never be enough time to do it all. Life is full of surprises, and sometimes an opportunity comes along that is truly too good to pass up, even if it means temporarily working at full capacity. "Full capacity" is itself a stretchable

concept for humans; working a few extra hours or temporarily hiring an assistant during a peak can make a big difference in keeping to-do lists and response times under control, but that should be the exception; as a norm, it is not sustainable. “Slack” does not have to mean doing nothing at all; spending time on preventive maintenance, recharging physically or emotionally, is a perfectly legitimate use of slack. A former business development manager realized that “maintaining deliberate slack is not laziness but a rational strategy.”

Lever 2: reduce variability

You can reduce variability by scheduling recurring tasks ahead of time, proactively smoothing the incoming flow. This is the idea behind metering on-ramps on Californian freeways: by reducing the variability of merging traffic, the freeway flows more smoothly. If April is always busy because of tax season, schedule your medical check-ups for a different part of the year. Quarterly reports are a classic and often poorly managed cause of spikes in demand. The other variability is in how long things take. Using routines, standard procedures, and checklists help to reduce this variability. This may require some upfront investment in documenting processes, which can make future work both faster (lower utilization) and more predictable (lower variability). Creating standard process flowcharts and templates for recurring reports can make a big difference. A former manager at a major bank put it as follows: “operations teaches that the goal isn’t maximum busyness, it’s smooth flow.” Variability is a tricky concept to grasp and not immediately visible, so often neglected; the analytical lens of queueing theory shows how impactful it can be, for better or worse. The only way to be both responsive and busy is by controlling variability.

Walk the Toyota Way

We have argued that you are your bottleneck, and that you need slack to accommodate the inevitable variability. Slack does not have to mean goofing off; instead, that slack can be

reinvested in small changes, that then lead to further improvements, and so on. Slack is what enables continuous improvement. What do those small changes look like? In the language of *The Goal*, how do we keep optimizing and elevating the bottleneck? The Toyota Way, the gold standard of self-propelling continuous improvement, is not about machines and robots; first and foremost, it is about people. Toyota itself says "It is based on the premise of making work easier for workers."[17] Although Toyota is known for making cars, the underlying principles of the Toyota Way have been applied in many other settings, including software development, health care, and non-profits.

Slack is what enables continuous improvement.

A core part of the Toyota Way is to avoid waste (or "muda") of effort, talent, time, money, resources, space, pollution, whatever. To do that, the Toyota Way prescribes identifying and fixing the root causes of waste, rather than hiding them. This is hard: our human tendency is to want to get stuff done quickly, rather than do the hard work of asking why something is happening. What does continuous improvement look like in the Toyota Way, and how does that map into common time management practices? We do not propose that you should adopt all of the practices that follow; find what works, and keep experimenting and improving. The Toyota Way is not a bag of tools, but a philosophy, so take what follows in that spirit.

To-do lists and kanban boards: make it visual

Time management gurus invariably preach creating clear visual lists. This helps you know what to do, and forces you to organize tasks according to your priorities. It also helps by externalizing your memory, which reduces cognitive load: once you put everything on a to-do list, you don't need to worry anymore about forgetting something. A health care analyst told us "getting things out of my head and into a system makes it easier to focus and be present." To-do lists are, paradoxically, very liberating. An experienced consultant shared "I started implementing it this past week and already feel a noticeable sense of clarity and relief once those items are out of

my head and organized somewhere else." The list provides a visual warning when you are taking on too much at once, making it easier to say "no." Lists are shareable, making it easier to coordinate and delegate.

The Toyota Way is also very visual. Workers easily see what they need to work on. Sometimes a "kanban" board shows where each job is in the process, so each worker knows exactly what to work on next. Kanban boards have become common in agile software development (itself influenced by the Toyota Way), where they help teams visually keep track of projects. Tools such as Jira and Asana include kanban boards for teams, and a variety of tools exist for individuals too. For others it can be as simple as a collection of post-it notes on a whiteboard. A software engineer described how they limit work-in-process and only "pull" the next task once the previous one is finished.

Build routines; standardize work

Doing something for the first time is often difficult. It takes time and willpower to start a new practice, whether making to-do lists or exercise, even if you know it's good for you. Time management recommends routines: habits replace willpower and conserve cognitive effort. (Recall Tim Cook's Apple cafeteria routine.) Some authors recommend planning your outfits for the entire week ahead of time, so you don't spend time and energy in the morning thinking about what to wear. A portfolio manager commented that "operations theory reframes time management as system design rather than willpower."

Doing the same task the same way every time reduces variability in processing times, workload and quality. This sounds boring, but standardization is a requirement for creativity: it is impossible to genuinely improve a process that is all over the place. Standardization does not mean never making changes, but eliminating unnecessary variability. The Toyota Way does not like waste, and unnecessary variability leads to wasted effort.

Foolproofing; poka-yoke

Spending bottleneck time to re-do work that was not done correctly first time is a big no-no, which is why the Toyota Way aims to make processes foolproof, known as "poka-yoke." Many practices become more powerful if they are made foolproof. If you need periods of uninterrupted time to work, you can physically unplug your internet router to block the temptation to check your email. If that is too drastic, a variety of tools exist to block distractions on your phone or laptop. Putting your phone charger in a fixed location away from your workplace keeps your phone (with all its distractions) out of sight when you need to do deep work, and you always know where it is. Some email programs display a pop-up message "you may have forgotten an attachment; send anyway?", another example of foolproofing. Both Gmail and Outlook let you delay outgoing email to allow you to catch mistakes, such as sending an email to the wrong people, or sending something you will later regret and need to spend time putting right. A finance manager at a major entertainment company observed that "instead of buffering in time for us to triple-check everything, we should dummy-proof our models and prep our decks more."

Inbox zero – control Work-in-Process (WIP)

Email is a top source of stress for many overworked professionals. How does inbox zero help? Every time you open your inbox, your mind subconsciously processes every email, even those it has seen before, wasting cognitive resources, another of Levitin's observations.[4] Instead, when you open your inbox, you should immediately take one of three actions: delete it, deal with it immediately (if it can be done quickly enough), or put it on your to-do list and perhaps schedule a future time to process it. Often enough, an email you had scheduled for later will have been resolved by then. Inbox zero may sound daunting, but people tend to find it incredibly liberating; just imagine email no longer being a source of stress in your life. The improved search capabilities of email software do not make inbox zero redundant, as you will still waste cognitive capacity looking at all those emails, divining which ones need attention, even if you

have flagged others as read—or prompting an AI agent to do so on your behalf. Inbox zero does not need to mean “zero at all times”; there will be inevitable spikes, but having a system to prevent the spikes from getting out of hand and regularly bringing the inbox back to zero is what matters.

Inbox zero is closely related to several core practices in the Toyota Way, such as Just-In-Time and the 5S framework for a clean and organized workspace (often translated as “Sort, Straighten, Shine, Standardize, Sustain”). Large amounts of work-in-process (WIP) cause delays, which means mistakes get found later. High WIP makes it hard to know what to work on, as everything eventually becomes urgent. High WIP causes errors and makes it easy to hide them, until they come back and bite. A fintech manager wrote “when my desk is cluttered, my inbox is full of unread emails, and I have dozens of browser tabs open, it becomes hard to see what actually matters. By keeping my workspace clean and my to-do's list lean, I can identify problems earlier and focus on what truly needs attention.”

Fix annoyances at root; pull the andon cord and employ kaizen

In the Toyota Way, whenever a worker encounters a problem, they pull the “andon” cord, a way of signaling that something needs attention and possibly stopping the production line temporarily to deal with it. We need the same ability to stop what we’re doing and step back. If we’re in the middle of something frustrating, we should ask ourselves whether there isn’t a better way or why we’re doing this at all. You open your laptop, energized for a productive session, but the software was updated overnight and is not behaving the way you are used to; some default settings (such as where to save files) have changed and now you have to manually keep modifying them. It’s annoying but you’re too busy to fix it, so you just keep spending time (and energy) fighting with the software. The Toyota Way would tell you to stop, identify the root cause of what is happening, and do a quick web search to find how to reset the defaults to what you want. When you are genuinely too busy to fix this now, schedule it for later: don’t let these

annoyances accumulate. Any one of them by itself is minor, but over time they can seriously hinder your productivity. You should control the software, not the other way around.

We can apply continuous improvement, or “kaizen”, to anything. We are all still learning how AI can work for us. It is easy to give it a prompt, such as “take this report and write a 1-page executive summary for my boss” and then be disappointed with the output and give up on AI for that purpose altogether. Instead, we should ask ourselves why it didn’t work, experiment with better prompts (AI can even help with that), and keep improving how we use AI.

Parkinson’s Law; eliminate waste (muda)

Parkinson’s Law is the principle that work expands to fill the available time; time management gurus therefore advise to not allocate too much time for any given task. If you think writing a memo will take four hours, try scheduling it for three hours. Pragmatism is important: if it’s really not finished in three hours, don’t send it to your boss; even if it takes three and a half hours, that’s thirty minutes saved. How does this fit with our earlier mantra to “create slack”? Simple: you create slack in your day by eliminating (unnecessary) slack in individual tasks.

Reducing waste, or muda, is at the very core of the Toyota Way. This is where the notion of “lean” operations comes from. A system with too much slack feels no pressure to improve. If your deadlines are too relaxed, unnecessary recurring delays will not bother you, so you feel no need to eliminate them. Maybe every project gets delayed towards the end because the customer’s requirements were not clear, and you end up making last-minute adjustments. Even if you do not miss the deadline, you are still wasting time and energy. Setting more ambitious deadlines (reducing slack) would force those problems into the open, nudging you to find a way to clarify customer expectations earlier, to avoid the back-and-forth at the end.

Operate with slack; avoid overburden (muri)

We have argued both for and against slack. In *Factory Physics*, Hopp and Spearman articulate the OM perspective on slack well.[18] Some slack (which they refer to as "buffer") is essential to accommodate the variability that you have not (yet) been able to get rid of, and that slack can be in the form of inventory, capacity, or time. For a writer, inventory can be an inventory of posts ready to go on Substack, capacity is how many posts they can produce per week, and time refers to the interval until the next post goes live. If the time between posts doesn't matter, or if capacity is much higher than the work involved, then no inventory is needed. How much slack is needed, and what type, depends on the variability one faces. The Toyota Way requires slack, just not unnecessary slack. It explicitly argues against overburdening workers (muri), as excessive workload is unsustainable and makes it impossible to invest in continuous improvement. Without some slack, no time management practice is going to work. They almost always require an investment of some time upfront to save more time later. Slack in resources is what creates the snowball effect of small changes continuously paving the way for further improvements (but too much slack in tasks stops the snowball in its track). A marketing executive was grateful that "knowing that slack can be healthy and a crucial component for achieving continuous improvement makes me feel less guilty."

The Toyota Way requires slack,

just not unnecessary slack.

Three closing thoughts on the Toyota Way and you

For the Toyota Way to work in practice, there are several important preconditions.

First, employees who do not experience psychological safety cannot follow the Toyota Way. A worker afraid of being yelled at will not signal a defect or suggest tighter deadlines. The same applies to you. Accept that not everything will work and don't beat yourself up about it, but you don't know until you try.

"Trying" is the second prerequisite: a culture of experimentation. The Toyota Way emphasizes the importance of "gemba" or "genchi genbutsu", going to see how the work actually gets done, in order to understand how to improve it. Experimentation in time management means trying out new practices in your own daily workflow. You are unlikely to immediately settle on the exact form of to-do lists that works best for you, but you can keep iterating to something better.

The third prescription is to be pragmatic. The Toyota Way does not aim to solve all problems at once, nor can you. Toyota is not perfect, and never will be; neither will we. All three of us are in the early stages of this journey too. Don't aim for the stars. Start making some small changes here and there, see what works and build on those; adapt or discard those that don't work.

When is the time to start managing time?

We have read many time management books, papers, and blogs, and have been struck how many of their recommendations map perfectly into three core pillars of operations management: managing by the bottleneck, queueing theory, and the Toyota Way. Table 1 provides an overview of how these OM principles underlie common time management practices. What Table 1 does not convey well is why many of these scattered time management practices are more powerful when viewed cumulatively rather than in isolation. A small change by itself creates the conditions that allow for further small changes, so the effects compound over time. The time management literature undervalues this. We read some advice, like inbox zero, and think "that sounds fine but surely that can't make enough difference to be worth doing." It *is* worth doing: inbox zero is a revelation for most people once they get there, but its real power comes from all the downstream changes it enables.

The practices in Table 1 are all closely connected to the principles of managing by the bottleneck and to the consequences of queueing theory. Consider the recommendation to outsource (in step 2 of managing by the bottleneck): it is much easier to identify the top

candidates for outsourcing if you have a list of tasks in front of you. Outsourcing will likely require an initial investment of time to describe the task in enough detail that someone else can do it; that process of documenting it will help standardize the task, and reduce variability in processing time and quality. Having that to-do list will allow you to schedule tasks ahead of time, including scheduling blocks of time for deep work; those longer uninterrupted blocks will allow you to get more done (optimizing the bottleneck), and fewer interruptions means less variability and probably fewer errors (hence contributing to a shorter to-do list). If we tried to visually illustrate how our recommendations about bottlenecks, queueing, and continuous improvement relate to one another, it would be a spaghetti diagram; all of these practices and principles are related to each other. That is why it is so important to view managing by the bottleneck and the Toyota Way as philosophies rather than toolkits, and queueing theory as a mathematical prescription linking circumstances to outcomes.

It doesn't matter where you start; the point is to start somewhere, and the busier you feel, the more valuable it is to start doing something about that. The text box gives an example of how, without intervention, busy turns to burnout, and how small improvements in time management can turn that vicious cycle into a virtuous one.

You may still be thinking, "that all sounds great, but I don't have time right now, I'll do it later." We have explored this question theoretically,[19] but the fable of the tortoise and the hare gives the same answer in a more relatable way: the best time to start is now. Why? Simple: the benefits compound. It's a true snowball effect: the longer your snowball rolls, the bigger and more self-propelling it becomes. The person who starts rolling a snowball at the top of the hill will soon have a juggernaut, while the person who walks halfway down the hill and only starts rolling then will never catch up. When it comes to money, we understand the power of compounding very well, but we don't tend to apply that same logic to time. Consider the following: most people would gladly invest $5,000 now to earn $1,000 per month forever, but far fewer will spend five hours now to fix a problem costing them one hour every month. Their time discount rate is far higher than their money discount rate.[20]

You arrive at work already feeling behind. New requests, emails, meetings, and unexpected problems continue to arrive throughout the day, making it impossible to follow the plan you had in mind that morning. As unfinished work accumulates, deciding what deserves attention becomes increasingly difficult, and tasks take longer than they should because your attention is scattered across too many priorities. The growing backlog creates stress and constant urgency, which further slows progress and leaves you falling even farther behind. Before long, simply keeping up consumes all your available time, leaving no opportunity to address the underlying causes of the problem.

Now imagine a different approach. You begin by saying no more often, limiting the number of projects in progress, maintaining a visual task list, and reserving a small amount of buffer time each week. With fewer items competing for attention, your inbox shrinks, multitasking declines, and recurring work becomes easier to schedule, template, and delegate. The time freed by these changes is reinvested in further improvements, such as documenting procedures, eliminating recurring annoyances, strengthening routines, and protecting periods of deep work. What started as a handful of small practices evolves into a virtuous cycle: better management of your personal bottleneck reduces variability, lower variability creates slack, and that slack enables the continuous improvements that make every subsequent hour more productive than the last.

Box: illustration of the compounding effect of time management practices

Most people would gladly invest $5,000 now to earn $1,000 per month forever, but far fewer will spend five hours now to fix a problem costing them one hour every month.

What will your time-money exchange rate look like in future? If your career is advancing, or your business is growing, or the value of your time is increasing for any other reason such as expanding family obligations, the opportunity cost of those five hours in future will be even higher, which will make it even harder to invest that time later. The longer you wait with investing in time management, the harder it gets. There is no time like the present to present yourself with more time in the future.

Author bios



Charles J. Corbett is IBM Chair of Management and Professor of Operations Management and Sustainability at the UCLA Anderson School of Management, where he has served as deputy dean and vice chair. He is a (Distinguished) Fellow of the Production and Operations Management Society and the Manufacturing and Service Operations Management Society. His current work focuses on the operations of well-being.

Email: charles.corbett@anderson.ucla.edu.



Guillaume Roels is the Timken Chaired Professor of Global Technology and Innovation and the Academic Director of the Global Executive MBA program at INSEAD. His research lies on the interface of operational excellence, people-centric operations, and the management of services.
He is currently serving as the Editor-in-Chief of *Service Science* and as a board member of the INFORMS Service Science section and the POMS College of Service Operations.

Guillaume.roels@insead.edu



Onesun Steve Yoo is Professor of Operations and Marketing Analytics at UCL School of Management and a co-director of the UCL Centre for Sustainable Business, University College London. His research examines how time shapes innovation and entrepreneurship. He currently serves as a senior editor of *Production and Operations Management* and an associate editor of *Manufacturing and Service Operations Management.*

Email: onesun.yoo@ucl.ac.uk

Acknowledgments

We are grateful to two anonymous reviewers for their insightful feedback on an earlier version of this manuscript, and to the Editor-in-Chief, Professor Kalyan Singhal, for his encouraging feedback.

Endnotes

1. Goldratt, E.M., Cox, J. (1984). *The Goal: A Process of Ongoing Improvement*. North River Press.
2. Liker, J.K. (2004) *The Toyota Way, 14 Management Principles from the World's Greatest Manufacturer*. McGraw-Hill, New York.
3. The quotes throughout this article are paraphrased from MBA students we have taught over the years, and are included to illustrate how framing time management through an operations management perspective has worked for them.
4. Levitin, D.J. (2014). *The Organized Mind: Thinking Straight in the Age of Information Overload*. Dutton.
5. https://fortune.com/2025/01/24/apple-ceo-tim-cook-strict-routine-eat-the-same-cafeteria-every-day-driving-factor-behind-his-success/, last accessed November 4, 2026.
6. Salvucci D.D., Taatgen, N.A. (2011). *The Multitasking Mind*. Oxford University Press, New York.
7. Newport, C. (2016). *Deep Work: Rules for Focused Success in a Distracted World*. Grand Central Publishing.
8. Holmes, C. (2022). *Happier Hour: How to Beat Distraction, Expand Your Time, and Focus on What Matters Most*. Simon and Schuster.
9. Whillans, A. (2020). *Time Smart: How to Reclaim Your Time and Live a Happier Life*. Harvard Business Review Press.
10. Drucker, P.F. (1967). *The Effective Executive*. Harper & Row.

11. Porter, M.E., Nohria, N. (2018). How CEOs manage time. *Harvard Business Review*, July-August, 2-11.
12. Newport, C. (2024). *Slow Productivity: The Lost Art of Accomplishment Without Burnout*. Portfolio.
13. Kim, G., Behr, K., Spafford, G. (2013). *The Phoenix Project: A Novel about IT, DevOps, and Helping Your Business Win*. IT Revolution.
14. Yoo, O.S. (2025). *Master Your Timeflow: Modern Guide to Time Management*. Working Insights Ltd.
15. Burkeman, O. (2021). *Four Thousand Weeks: Time Management for Mortals*. Farrar, Straus and Giroux.
16. Allen, D. (2015). *Getting Things Done: The Art of Stress-Free Productivity (Revised edition).* Penguin Books.
17. https://global.toyota/en/company/vision-and-philosophy/production-system/; last accessed December 27, 2025.
18. Hopp, W.J., Spearman, M.L. (2008). *Factory Physics* (3rd edition). Waveland Press.
19. Yoo, O.S., Corbett, C.J., Roels, G. (2016). Optimal time allocation for process improvement for growth-focused entrepreneurs. *Manufacturing & Service Operations Management*, 18(3), 361-375.
20. Zauberman, G., Lynch Jr, J.G. (2005). Resource slack and propensity to discount delayed investments of time versus money. *Journal of Experimental Psychology: General*, 134(1), 23-37.

Table 1: mapping selected time management practices to Operations Management principles

Time Management Practice	Related OM Concept	Why It Works (mechanism)
Build routines / habits	Standardize work	Reduces variability in processing time; frees cognitive capacity
Maintain a to-do list	Kanban / visual management	Prevents bottleneck idling; externalizes memory; reveals current workload, makes it easier to say no
Deep work; avoid multitasking	Batching; reduce setup cost	Task-switching wastes bottleneck time; larger batches amortize fixed mental setup cost
Outsource	Reduce workload	Reduces arrivals; creates slack at bottleneck; fewer items in queue
Say no / decline tasks	Reduce workload	Reduces arrivals; creates slack at bottleneck; fewer items in queue
Track how you use your time	Measure the opportunity cost of the bottleneck	Makes opportunity cost visible; surfaces hidden waste
Delegate	Outsource non-bottleneck work	Opportunity cost logic: your bottleneck time freed; assistant's time is cheaper
Ensure all info before starting	Move QC in front of bottleneck	Eliminates rework, which is doubly costly at the bottleneck
Don’t wait for non-bottlenecks	Avoid wasting bottleneck time	By definition non-bottleneck steps have slack, so they should wait for you, not vice versa
Avoid pseudo-productivity	Reduce workload	Keeping non-bottleneck steps busy typically creates more work for the bottleneck too
Exercise, diet, sleep, mindfulness	Preventive maintenance	Prevent breakdowns which lead to lost bottleneck time
Create unscheduled buffer time	Operate with slack	Too little slack => high WIP. Slack allows / is essential for continuous improvement.
Schedule recurring tasks in advance	Reduce demand variability	Smooths arrival rate => less WIP
Use templates; AI; document procedures	Reduce service variability	Less variance in task duration => less WIP
Foolproofing	Poka-yoke	Eliminates rework; reduces variability; makes good habits automatic and durable
Inbox zero	Control WIP	Eliminates repeated cognitive re-processing of same emails
Take on fewer projects at once	Control WIP	Lower WIP => faster turnaround, completion; fewer projects to keep track of, less switching
Fix annoyances at root; don't patch	Andon cord; kaizen	Allowing annoyances to continue has large hidden cost; saying “no time to fix it now” ignores that you have no time later either
Parkinson’s Law: set tighter deadlines (where possible)	Eliminate waste (muda)	Slack in individual tasks masks inefficiency and dulls improvement incentives
Start better time management now	Continuous improvement: start small, start now	As one constraint is lifted another emerges; improvement is a loop, not a destination

Note: many common time management practices are consistent with well-established OM principles related to bottlenecks, queueing theory, and the Toyota Way. The practices are listed here in the sequence in which they are mentioned in the text.