Stories at Work

By Robert Akerlof, Niko Matouschek, and Luis Rayo*

Chiat/Day, the iconic ad agency behind Apple's "1984" and "Think Different" campaigns, was driven by a rebel spirit. In their view, the Madison Avenue establishment were square suits who shut down creative voices and pandered to clients. By contrast, Chiat/Day with its surfer culture, loose dress code, and Venice Beach headquarters shaped like a giant pair of binoculars, took pride in breaking the rules and preaching rebellion. According to founder Jay Chiat: "We're the pirates, not the navy."[1]

Stories, like the one told at Chiat/Day, are commonplace. They are a crucial force shaping employee behavior: they affect knowledge and beliefs (e.g., Bénabou, Falk and Tirole, 2018, Gibbons and Prusak, 2020), serve as "mental models" (e.g., Cremer, Garicano and Prat, 2007, Mullainathan, Schwartzstein and Shleifer, 2008, and Gibbons, LiCalzi and Warglien, 2018), and directly influence preferences (e.g., Akerlof and Kranton, 2005). Presumably for this reason, organizations spend considerable resources constructing and disseminating stories to their workforces.

Organizations display a variety of stories. The Mayo Clinic, for instance, emphasizes teamwork and patient-first care. In the words of William J. Mayo: "The best interest of the patient is the only interest to be considered...[Consequently, it] has become necessary to develop medicine as a cooperative science."[2] Lincoln Electric, in contrast, has a philosophy of rugged individualism. According to James Lincoln, "Competition will mean the disappearance of the lazy and incompetent."[3] A different story prevailed at Arthur Andersen, where accountants prided themselves on being "androids" who took a traditional approach to business and scrupulously followed long-established procedures. The firm's ideals were manifested in its logo: the sturdy mahogany double doors of its Chicago headquarters.

Stories are shaped by a combination of top-down and bottom-up forces. On the top-down side, organizations disseminate stories that are advantageous for recruiting and motivating workers. For example, the Mayo Clinic, recognizing the importance of maintaining its culture, hires mostly from its own medical school where the training process is referred to as "Mayoizing." Likewise, Arthur Andersen put new recruits through centralized training at "Andersen U." On the bottom-up side, employees are more favorably disposed toward some stories than others. An employee's particular disposition is partly a matter of the prevailing cultural context. For instance, workers in more collectivist countries may be more willing to adopt a team-based story. Employees also have greater receptivity to stories that are identity enhancing (e.g., Guadalupe, Kinias and Schloderer, 2020). Chiat/Day's story, for instance, complemented the identities of its employees, many of whom did not fit easily into the traditionalist culture of Madison Avenue.

Here, our focus will be on top-down design and on a firm's use of stories, in particular, to motivate workers. We build a model where an organization must find a way to motivate employees both to work

* Akerlof: Department of Economics, University of Warwick, Coventry CV4 7AL, UK, email: r.akerlof@warwick.ac.uk; Matouschek: Kellogg School of Management, Northwestern University, 2211 Campus Drive, Evanston, IL 60208, email: n-matouschek@kellogg.northwestern.edu; Rayo: Kellogg School of Management, Northwestern University, 2211 Campus Drive, Evanston, IL 60208, email: luis.rayo@kellogg.northwestern.edu. We are grateful to Theo Durandard, Rob Gertner, Richard Holden, and Mike Powell for helpful comments.

[1] Tungate (2007), p. 113.

[2] Berry and Seltman (2008), p. 22.

[3] Berg and Fast (1975), p. 4.

hard and to work on the right thing, all the while keeping its budget balanced. The organization gets to choose the story of its employees and, in conjunction, their monetary incentives. As in Becker and Murphy (1993)'s model, we treat stories as a form of advertising—in our case to employees—that affects utility and influences decision-making.[4]

Our model predicts that an organization will adopt one of two designs. The first design, which we call a "purpose-driven" organization, pairs flat monetary incentives with a story that emphasizes the importance of generating output (e.g., saving lives, putting a person on the moon). The second one, which we call an "incentive-driven" organization, pairs steep monetary incentives with a narrower story that emphasizes the importance of maintaining ethical standards (e.g., maintaining quality, helping peers). Which design is optimal depends on the severity of multitasking and the relative costs of conveying different types of stories. In section III, we shall illustrate the applicability of these results.[5]

Our model is a first attempt at showing how an organization's choice of story from a rich set can be formally incorporated into its design problem. In our view, stories are key for understanding differences in organizational practices and, even more crucially, differences in firm productivity.

I. Baseline model

We begin with a simple baseline model without stories. This model builds on Holmstrom (1982)'s model of team production and Holmstrom and Milgrom (1991)'s model of multi-tasking. There is a team of $n \geq 2$ identical players. Player $i = 1, ..., n$ chooses two actions $a_i, b_i \geq 0$ at cost $c(a_i + b_i)$. All choices are simultaneous. These actions result in team output $y := f(\mathbf{a}, \mathbf{b})$ and individual performance measures $\phi_i := g(a_i, b_i; \mathbf{a}_{-i}, \mathbf{b}_{-i})$, where variables in bold represent vectors. Players have a symmetric impact on output and on the performance measures of their peers—and so, for instance, the first-best actions are the same for all players.[6] Let $\Delta\phi_i := \phi_i - \sum_{j=1}^{n} \phi_j/n$ denote how well a player performed relative to the team. Both the team's output and the players' performance measures are contractible, but their actions are not. We assume that c, f, and g are smooth functions with c strictly increasing and strictly convex, and f strictly increasing and concave. Moreover, we assume that the relative performance measure $\Delta\phi_i$ is concave in player i's actions, satisfies $\frac{\partial \Delta\phi_i}{\partial a_i} > 0$, and is more responsive to a_i than to b_i, that is $\frac{\partial \Delta\phi_i}{\partial a_i} > \left|\frac{\partial \Delta\phi_i}{\partial b_i}\right|$.

One interpretation is that a_i and b_i are two dimensions of effort devoted to a given task, such as "quantity" and "quality." Another interpretation is that a_i is effort devoted to an individual task, ϕ_i measures the player's performance in this task, and b_i is effort devoted to helping other players perform their own tasks. In this case, $\frac{\partial \phi_i}{\partial b_i} = 0$ and $\frac{\partial \Delta\phi_i}{\partial b_i} < 0$.

Player i's total compensation is $w_i(y, \phi)$, with $\sum_i w_i(y, \phi) = y$ for all y and ϕ, so that the team's budget is balanced. Player i's payoff is $w_i(y, \phi) - c(a_i + b_i)$. For simplicity we assume that $w_i(y, \phi)$ is linear and symmetric across players so that $w_i(y, \phi) \equiv y/n + \lambda\Delta\phi_i$ for some $\lambda \geq 0$, where λ is a constant that captures the strength of the players' monetary incentives.

The team's problem is to choose λ and prescribe actions so as to maximize joint surplus $f(\mathbf{a}, \mathbf{b}) - \sum_i c(a_i + b_i)$ subject to the constraint that a_i, b_i maximizes player i's payoff given $\mathbf{a}_{-i}, \mathbf{b}_{-i}$. Let $(\mathbf{a}^*, \mathbf{b}^*) := \arg\max_{\mathbf{a},\mathbf{b}} f(\mathbf{a}, \mathbf{b}) - \sum_i c(a_i + b_i)$ denote

[4] Our model differs from Becker and Murphy (1993) in two key respects. First, we consider a richer set of possible stories. Second, whereas Becker and Murphy (1993) consider the problem of motivating consumers, we consider a more complex problem of motivating employees in a multitasking environment.

[5] Our paper is related to Prendergast (2007) and Besley and Ghatak (2005) where, through selection, a firm chooses the intrinsic motivation of its workers. This choice is akin to the choice of story in our model. What distinguishes our approach is that stories come in a variety of flavors.

[6] Rayo (2007) also allows for individual performance measures in a team environment, but there is no helping or multitasking; Itoh (1991) allows agents to help one another, but these agents contract with a principal rather than with each other.

the first-best actions (which we assume are unique and interior), and let $y^* := f(\mathbf{a}^*, \mathbf{b}^*)$ denote first-best output.

Crucially, because $\frac{\partial \Delta\phi_i}{\partial a_i} > \frac{\partial \Delta\phi_i}{\partial b_i}$, the team faces a non-trivial multi-tasking problem and the first-best actions cannot be attained. The reason is standard: because output must be shared across players, the team needs to rely, at least in part, on individual incentives λ; but since the performance measure is biased toward a_i, these incentives inevitably cause a distortion.

II. Stories

We now bring stories into the picture. We assume that the team chooses a story, common to all players, that serves as an additional source of incentives. This story will allow the team to implement the first best.

Player i's payoff is now $w_i + s_i - c(a_i + b_i)$, where the new term s_i represents the player's *story utility.* We assume that s_i depends upon the player's actions and team output and takes the following form:

$$s_i(a_i, b_i, y) = \alpha(a_i - a^*) + \beta(b_i - b^*) + \tau(y - y^*)$$

The particular story chosen by the team is represented by non-negative constants α, β, and τ. A story with a positive α is one that emphasizes the importance of the first action (e.g., producing a high volume of output); a story with a positive β emphasizes the importance of the second action (e.g., producing high quality output, or helping peers); and a story with a positive τ emphasizes the value of team output per se (e.g., taking pride in the company's mission).

We have assumed that players compare their actions as well as the team's output against the first-best levels. This assumption serves two purposes. First, it allows us to focus on the incentive effect of stories, rather than their impact on utility (in equilibrium, story utility will be zero). Second, it adds realism to the model as there is considerable evidence that people engage in comparisons (see Akerlof, 2017).

While the team gets to choose any story it wishes, it must devote resources to constructing it and disseminating it to its members. We assume, in particular, that the cost of the story is $\alpha + \beta + T \cdot \tau$ for some parameter $T > 0$; that is, the firm faces constant returns to scale in the intensity of the story, and for simplicity the marginal costs of α and β are equal (and normalized without loss to 1).[7]

We shall focus on a special but instructive problem: choose the cheapest possible combination of monetary incentives λ and story α, β, τ subject to the constraint that this combination induces players to choose the first best.[8] Dropping the i subscripts for α_i, b_i, and $\Delta\phi_i$, the team's problem is:

$$\min_{\lambda,\alpha,\beta,\tau \geq 0} \ \alpha + \beta + T \cdot \tau \tag{1}$$

$$s.t.$$

$$(1/n + \tau) MC^* + \lambda \frac{\partial \Delta\phi^*}{\partial a} + \alpha = MC^*,$$

$$(1/n + \tau) MC^* + \lambda \frac{\partial \Delta\phi^*}{\partial b} + \beta = MC^*,$$

where $MC^* := c'(a^* + b^*) = \frac{\partial f(\mathbf{a}^*, \mathbf{b}^*)}{\partial a} = \frac{\partial f(\mathbf{a}^*, \mathbf{b}^*)}{\partial b}$, $\frac{\partial \Delta\phi^*}{\partial a} := \frac{\partial \Delta\phi(\mathbf{a}^*, \mathbf{b}^*)}{\partial a}$, and $\frac{\partial \Delta\phi^*}{\partial b} := \frac{\partial \Delta\phi(\mathbf{a}^*, \mathbf{b}^*)}{\partial b}$. The objective is the cost of the story and the two constraints are the players' first-order conditions for a_i and b_i, respectively, evaluated at the first best. The left-hand side of each equation measures the marginal benefit of a higher action, and the right-hand side measures its marginal cost. Because of our assumptions, these conditions are sufficient for implementing the first best.

We now proceed to solve the team's problem. We refer to a solution as an *optimal design*, and to the corresponding story as an *optimal story.* In what follows, we focus on environments that are "generic" in the sense that the linear program defined by (1) has a single solution.[9] The only instances where this in not the case are knife-edge.

[7] By assuming that α and β have equal marginal costs, we avoid biasing the story in favor of one action over the other.

[8] In practice the team may opt for a less expensive story that does not implement the first best. We conjecture that so long as the desired actions remain close to first best, the optimal story will remain close to the one we derive here.

[9] We assume, in particular, that $T \neq MC^*(1 - \frac{\partial \Delta\phi^*}{\partial b} / \frac{\partial \Delta\phi^*}{\partial a})$.

LEMMA 1: *Every optimal story has* $\alpha = 0$.

To see why this is true, begin with a story that motivates the first action a, at least in part, through a positive α. The same action a can be obtained by raising monetary incentives λ and lowering α. Because these stronger monetary incentives have the potential to aggravate the multitasking problem, β might need to be raised. Even so, since the contractible performance measure $\Delta\phi$ is more responsive to a than to b, α would fall more than β would grow; hence, overall story costs would fall.

LEMMA 2: *Every optimal story has either* $\tau = 0$ *or* $\beta = 0$.[10]

This result tells us that a story that rewards players for total output ($\tau > 0$) substitutes for a story that rewards them, more narrowly, for action b ($\beta > 0$). The reason is that a τ-story spreads incentives evenly across both actions, therefore reducing the need for a β-story that deals with the players' multitasking problem. Moreover, because the story technology has constant returns, the team goes to an extreme where only one of these stories is employed.

THEOREM 1: *The optimal design takes one of the following two forms.*[11]

1. Purpose-driven organization:

$$\tau = \frac{n-1}{n} \text{ and } \lambda = \beta = 0.$$

2. Incentive-driven organization:

$$\lambda, \beta > 0 \text{ and } \tau = 0.$$[12]

The first design is optimal if and only if T, *the marginal cost of an output-based story, is below the threshold* $MC^* \left(1 - \frac{\partial\Delta\phi^*}{\partial b} / \frac{\partial\Delta\phi^*}{\partial a}\right)$, *which measures the severity of multitasking.*

This result tells us that the team will either adopt an output-based story and remove all monetary incentives, or will instead offer high-powered monetary incentives (sufficient to induce a first-best level of action a) together with a narrower story that rewards action b (sufficient to address multitasking). We interpret this narrower story as an ethical standard concerned with maintaining quality or helping peers. In other words, monetary incentives and τ-stories are substitutes, whereas monetary incentives and β-stories are complements. Intuitively, because a τ-story provides incentives for both actions, and these incentives are perfectly balanced, there is no need for any additional instrument. In contrast, monetary incentives can only achieve the first best when used in tandem with a narrower β-story that counteracts the multitasking bias.

Observe that, other things equal, the team will be purpose-driven when the cost of telling an output-based story is sufficiently low. Thus, firms that produce the types of output that employees naturally value will gravitate toward this design. The team will also be purpose driven when the multitasking problem is sufficiently pronounced,[13] for in this case monetary incentives become unattractive.

III. Applications

Here we present some examples of firms that are organized in a way broadly consis-

[10] To see why, suppose $\tau, \beta > 0$. Because $\alpha = 0$, we must have $\tau < (n-1)/n$ and $\lambda > 0$. Now choose a small $\varepsilon \neq 0$, change τ by ε/MC^*, change λ by $-\varepsilon/\frac{\partial\Delta\phi^*}{\partial a}$, and change β by $-\varepsilon(1 - \frac{\partial\Delta\phi^*}{\partial b}/\frac{\partial\Delta\phi^*}{\partial a})$. Both constraints are still met and story costs change by $\varepsilon(T/MC^* - 1 + \frac{\partial\Delta\phi^*}{\partial b}/\frac{\partial\Delta\phi^*}{\partial a})$. Generically, this change is either strictly positive or strictly negative, which is a contradiction.

[11] The proof of the theorem is as follows. When $\beta = 0$, the bias in the performance measure (i.e. $\frac{\partial\Delta\phi^*}{\partial b}/\frac{\partial\Delta\phi^*}{\partial a} < 1$) means that $\lambda = 0$. As a result, we must have $\tau = \frac{n-1}{n}$. In this case, the per-player cost of the story is $\frac{n-1}{n}T$. When instead $\tau = 0$, we must have $\lambda = \frac{n-1}{n}\frac{MC^*}{\partial\Delta\phi^*/\partial a}$ in order to motivate action a, and must have $\beta = \frac{n-1}{n}MC^*(1 - \frac{\partial\Delta\phi^*}{\partial b}/\frac{\partial\Delta\phi^*}{\partial a})$ in order to address multitasking. In this case, the per-player cost of the story is $\frac{n-1}{n}MC^*(1 - \frac{\partial\Delta\phi^*}{\partial b}/\frac{\partial\Delta\phi^*}{\partial a})$. The last part of the theorem follows from comparing the two story costs.

[12] Specifically, $\lambda = \frac{n-1}{n}\frac{MC^*}{\partial\Delta\phi^*/\partial a}$ and $\beta = \frac{n-1}{n}MC^*(1 - \frac{\partial\Delta\phi^*}{\partial b}/\frac{\partial\Delta\phi^*}{\partial a})$.

[13] This occurs when the first-best actions are sufficiently costly (so that MC^* is high), when the performance measure is sufficiently biased (so that $1 - \frac{\partial\Delta\phi^*}{\partial b}/\frac{\partial\Delta\phi^*}{\partial a}$ is large), or some combination of the two.

tent with our predictions.

The Mayo Clinic is an example of a purpose-driven organization. Its story emphasizing teamwork and patient-first care is a prime example of a τ story. The ease of telling this story is reflected in the anecdote of a custodian who, when asked why she was working so hard cleaning a hospital room, responded "I'm saving patients' lives."[14] In line with the model's predictions, Mayo complements its τ story with a flat compensation structure under which physicians' pay depends only on specialty and tenure. This structure was designed deliberately to encourage efficient referrals and cooperation (see Berry and Seltman, 2008).

Another powerful τ story can be found in the U.S. Marine Corps. Their story, built in part through a grueling boot camp, emphasizes team spirit, equality, and willingness to sacrifice oneself for the cause, as encapsulated in their maxim: "Every Marine is a rifleman." Using this story is not without costs. Recently, the Corps has faced difficulties filling positions in critical high-tech areas such as cyber security. They have considered offering "lateral entry" to civilians with the necessary expertise, starting them at mid-career rank and exempting them from boot camp. The marines are concerned, though, that such lateral hires will make it harder to maintain their story (in terms of our model, a higher cost T). Consequently, there is reticence about the practice. As retired Lt. Col. Dakota Wood points out, "If you go away from [the maxim], then I think you lose something that has made the Marine Corps what it is...If that breaks down, you've got problems."[15,16]

These examples contrast with professional service firms such as McKinsey, Goldman Sachs, and Capital Group. In line with the incentive-driven organizations predicted by our model, these firms combine high-powered monetary incentives with narrower stories extolling the virtue of "professionalism," the willingness to forgo short-term profits for clients' long-term needs, a form of β story. Capital Group, for instance, prides itself on going against the industry norm of "selling what sells," Goldman Sachs impresses on its employees that it would rather be "best than biggest," and McKinsey tells the story of how, at age 93, its iconic former managing director Marvin Bower interrupted a firm-wide meeting to remind partners that "If there is the shadow of a doubt on something being good for business but not truly professional, do not do it!" (see Ellis, 2013, pp. 4-16). Recent corporate scandals illustrate the dangers such firms are exposed to if their stories are not sufficiently compelling to induce employees to resist their monetary temptations.

Finally, to keep things simple, we have assumed that all players are identical and that they all share the same story. In practice, however, a firm might choose different stories for different parts of its workforce. Our model offers some hints as to how this might play out. Consider, for example, a team of homogeneous high-level employees (such as managers) who perform a broad set of tasks and thus face significant multitasking; separately, consider a team of homogeneous low-level employees (such as line workers) who perform more focused tasks where multitasking is minimal. Per Theorem 1, we expect the organization to spend significant resources in constructing a story for high-level employees (based on either τ or β), but not for low-level ones. Thus, the nature of the two jobs will be very different: high-level employees will experience, through their story, a sense of meaning in their work; whereas low-level employees will be rewarded primarily through monetary incentives (for example, through piece rate or efficiency wages) and will potentially, and unfortunately, find less meaning in their jobs. Concerning the role of work in peoples' lives, Derek Thompson (2019) notes: "for the poor and mid-

[14] Ellis (2013), p. 31.

[15] See Schogol (2017).

[16] An example of a for-profit firm with a purpose-driven organization is the international executive search firm Egon Zehnder. Like Mayo, its story emphasizes serving clients through teamwork and its pay structure is flat. Partners share profits equally to encourage information sharing and referrals of clients to whoever is in the best position to help them (See Zehnder, 2001).

dle class, work [is] a necessity; but for the college-educated elite, it [has become] a kind of religion, promising identity, transcendence, and community."

REFERENCES

Akerlof, George A, and Rachel E Kranton. 2005. "Identity and the Economics of Organizations." *Journal of Economic Perspectives*, 19(1): 9–32.

Akerlof, Robert. 2017. "Value Formation: The Role of Esteem." *Games and Economic Behavior*, 102: 1–19.

Becker, Gary S, and Kevin M Murphy. 1993. "A Simple Theory of Advertising as a Good or Bad." *The Quarterly Journal of Economics*, 108(4): 941–964.

Bénabou, Roland, Armin Falk, and Jean Tirole. 2018. "Narratives, Imperatives, and Moral Reasoning." *NBER Working Paper No. 24798*.

Berg, Norman, and Norman Fast. 1975. "The Lincoln Electric Company." *Harvard Business School Case 9-376*.

Berry, Leonard L, and Kent D Seltman. 2008. *Management Lessons from Mayo Clinic.* McGraw-Hill Professional Publishing.

Besley, Timothy, and Maitreesh Ghatak. 2005. "Competition and Incentives with Motivated Agents." *American Economic Review*, 95(3): 616–636.

Cremer, Jacques, Luis Garicano, and Andrea Prat. 2007. "Language and the Theory of the Firm." *The Quarterly Journal of Economics*, 122(1): 373–407.

Ellis, Charles D. 2013. *What it Takes: Seven Secrets of Success from the World's Greatest Professional Firms.* John Wiley & Sons.

Gibbons, Robert, and Laurence Prusak. 2020. "Stories as Knowledge in Organizations." *American Economic Review P & P*, Forthcoming.

Gibbons, Robert, Marco LiCalzi, and Massimo Warglien. 2018. "What situation is this? Coarse cognition and behavior over a space of games." *Mimeo*.

Guadalupe, Maria, Zoe Kinias, and Florian Schloderer. 2020. "Individual Identity and Organizational Identification: Evidence from a Field Experiment." *American Economic Review P & P*, Forthcoming.

Holmstrom, Bengt. 1982. "Moral Hazard in Teams." *The Bell Journal of Economics*, 324–340.

Holmstrom, Bengt, and Paul Milgrom. 1991. "Multitask Principal-Agent Analyses: Incentive Contracts, Asset Ownership, and Job Design." *Journal of Law, Economics, and Organization*, 7: 24.

Itoh, Hideshi. 1991. "Incentives to help in multi-agent situations." *Econometrica*, 59(3): 611–636.

Mullainathan, Sendhil, Joshua Schwartzstein, and Andrei Shleifer. 2008. "Coarse Thinking and Persuasion." *The Quarterly Journal of Economics*, 123(2): 577–619.

Prendergast, Canice. 2007. "The Motivation and Bias of Bureaucrats." *American Economic Review*, 97(1): 180–196.

Rayo, Luis. 2007. "Relational Incentives and Moral Hazard in Teams." *Review of Economic Studies*, 74(3): 937–963.

Schogol, Jeff. 2017. "Every Marine a rifleman no more?" *Marine Corps Times*, May 7.

Thompson, Derek. 2019. "Workism is Making Americans Miserable." *The Atlantic*, 24.

Tungate, Mark. 2007. *Adland: A Global History of Advertising.* Kogan Page Publishers.

Zehnder, Egon. 2001. "A Simpler Way to Pay." *Harvard Business Review*, 79(4): 53–6.