Organizational Change and Reference-Dependent Preferences†

Klaus M. Schmidt*
University of Munich, CESifo, and CEPR

Jonas von Wangenheim**
University of Bonn

This version: January 12, 2026

Abstract: Reference-dependent preferences can explain several puzzling observations about organizational change. We introduce a dynamic model in which a firm bargains with loss-averse workers about organizational change and wages. We show that change is often stagnant or slow for many periods, followed by a sudden boost in productivity triggered by a crisis. In a crisis, workers concede to organizational change but resist nominal wage reductions. In addition, the model explains why different firms in the same industry often have significant productivity differences. Finally, the model demonstrates the importance of expectations management even if all parties have rational expectations.

JEL classification numbers: D23; D91; L2.

Keywords: Organizational Change; Productivity; Reference Points; Loss Aversion.

---

†We would like to thank Christian Beyer, Thomas Dohmen, James Dow, Carl-Christian Groh, Paul Heidhues, Felix König, Botond Köszegi, Daniel Krähmer, Matthias Kräkel, Stephen Leider, Takeshi Murooka, Volker Nocke, Nicolas Schutz, Heidi Thysen, and seminar audiences in Bergen, Berlin, Bonn, Bristol, Hamburg, Heidelberg, Konstanz, Munich, Osaka, Tokyo, Zurich and the EARIE and EEA meetings for many helpful comments and suggestions. Financial support by Deutsche Forschungsgemeinschaft through CRC-TRR 190 (project number 280092119) and CRC-TR 224 (Project B03) is gratefully acknowledged.

*Department of Economics, University of Munich, Ludwigstr. 28 (Rgb.), D-80539 München, Germany, email: klaus.schmidt@lrz.uni-muenchen.de (corresponding author).

**Department of Economics, University of Bonn, Adenauerallee 24-42, D-53113 Bonn, Germany, email: jwangenheim@uni-bonn.de.

1 Introduction

In the early 1980s the iron ore and steel industry in the Great Lakes area was hit by a severe competitive shock. For more than 100 years it had been protected from foreign competition by its geographic location. Then it suddenly faced foreign competitors offering steel at substantially lower prices putting its own future into doubt. However, the local steel industry managed to survive and thrive by implementing cost reductions and boosting productivity. Within a few years, labor productivity doubled! This remarkable development was not due to the introduction of new technologies, nor to major capital investments, nor to the exit of low productivity firms. Instead, competitiveness was mainly restored by organizational improvements, in particular changes of inefficient work practices and the more efficient use of existing capital.[1]

This episode highlights three puzzles of organizational and technological change. First, despite substantial technological change, firms are often slow to adapt and to exploit new technological opportunities. However, change is sometimes implemented very rapidly, often in the course of a crisis.[2] Why does it need a crisis to adapt efficiently? Second, in a crisis workers have to make concessions to keep their firm in business. Firms use this to push through organizational changes, but rarely to cut nominal wages.[3] Why do wages not fall in a recession? Finally, different firms in the same industry facing the same technological opportunities often have large productivity differences for extended periods of time.[4] Why don't all firms move to the efficiency frontier?

In this paper, we show that reference-dependent preferences, in particular loss aversion, offer an explanation for these puzzles.[5] In our model technological change creates profitable

[1]Schmitz (2007) offers a detailed and thought-provoking case study of this episode.

[2]Holmes and Schmitz (2010) survey case studies documenting this pattern.

[3]Nominal wage cuts are sometimes observed, but they are rare, even in periods of low inflation and recessions. See Fallick, Villar, and Wascher (2020) and Branten, Lamo, and Room (2018).

[4]Syverson (2004) reports for the US that a firm at the 90th percentile of productivity has a total factor productivity (TFP) that is 1.9 times higher than the TFP of a firm at the 10th percentile (on average, in industries at the four-digit level). These differences cannot be accounted for by differences in observable inputs or heterogeneous prices. Increasing evidence suggests that differences in productivity are associated with organizational differences and differences in management practices. See e.g., Bloom et al. (2014), Bloom et al. (2019).

[5]Loss aversion is one of the most widely studied "biases" in behavioral economics. A recent meta-study by Brown et al. (2023) reports that the average parameter of loss aversion in 607 empirical estimates is

opportunities for a firm, but requires organizational changes. In the main part of the paper we abstract away from any informational or contractual frictions and assume that organizational change is perfectly contractible. However, implementing change is costly, because workers have to be compensated not only for the physical cost of adaptation, but also for the psychological loss associated with the change in their work practices. If there was no loss aversion the firm would implement the materially efficient change and compensate workers for the physical adaptation cost. With loss aversion, however, there is a threshold such that for all technological changes below this threshold the firm will stick to the status quo. If the threshold is exceeded the firm will adjust somewhat, but it will adjust less than in the situation without loss aversion. Thus, as is well known, loss aversion creates inertia.

This changes in a crisis. A crisis reduces the outside option of workers as the firm is threatened by bankruptcy, leading to job loss and lower paying jobs in other industries. In order to keep their jobs, workers have to make concessions, which can take the form of wage cuts or changes in work practices. Both of these concessions are perceived as losses by workers. To induce workers to accept organizational changes the firm can now offer a lower wage cut (a reduced loss), while it would have to offer a higher wage increase (a greater gain) in normal times. Thus, in a crisis the effect of loss aversion cancels out in the workers' marginal rate of substitution between wage payments and adaptation costs, and the firm will jump to the materially efficient organizational structure if there are sufficient worker rents to be appropriated. This offers a microfoundation for why wages do not fall in a recession: Workers make concessions by accepting change rather than by accepting wage cuts.

An important question for any model of loss aversion is how the reference point is determined. Kahneman and Tversky (1979) and much of the subsequent literature on loss aversion argue that in many situations a natural reference point is the status quo. In contrast, Köszegi and Rabin (2006) argue that the reference point should be determined by rational expectations given the decision maker's action. We take an agnostic position and assume that the reference point is some convex combination of the status quo and rational expectations. Our qualitative results hold for any combination that puts positive weight on both arguments.

This simple model captures the main intuition of why a crisis triggers organizational

$1 + \lambda = 1.955$, i.e., losses weigh almost twice as much as gains.

change, but it leaves several questions unanswered. The crisis has an effect only if workers enjoy a rent. Where does this rent come from? What happens if the reference point adjusts to the new contract and the new expectations? How do rational players prepare for the possibility of a crisis in the initial contract?

We develop a simple infinite horizon model that answers these questions. In this dynamic model there is technological change in every period. Loss aversion causes inertia, i.e., there is either no organizational adaptation to new technologies or less than material efficiency requires. If there is organizational change, workers suffer a one-period disutility due to adaptation cost and due to loss aversion associated with the adaptation. For these costs they are optimally compensated by a small but permanent wage increase. This gives rise to a quasi-rent that builds up over time. If the crisis hits, workers have to concede this quasi-rent by agreeing to substantial organizational change. If parties rationally anticipate the crisis, they will agree on higher wages to compensate for the expected utility loss in the crisis and they will delay change because change is cheaper to implement in the crisis. Thus, the model explains why there is often inertia or very slow change for extended periods of time, but then there is a sudden jump of productivity triggered by a crisis. Furthermore, the model explains why organizational change is history-dependent and may result in large productivity differences between firms that were founded at different points in time or faced idiosyncratic shocks.

So far, we assumed that change is both deterministic and perfectly contractible. In an extension of the model we consider a principal-agent problem between the owner of the firm and a manager who has to be incentivized to implement change stochastically. However, successful change always needs workers to go along with it. If the workers' reference point is (partially) determined by their expectations, an increase in the probability of successful change increases the reference point and makes it cheaper to pay workers to accept change. Thus, the owner will induce the manager to either implement the desired change with a high probability (or certainty) or to not implement it at all. This highlights the importance of expectations management, even if all parties form rational expectations. The finding is consistent with the emphasis on effective leadership, vision, and creating a sense of urgency for successful organizational change by practitioners and management consultants see, e.g, Kotter (1995) and Burke (2017).

Why do we need a model with loss aversion? Each of the phenomena described in the previous paragraphs can also be explained by other models. For example, a fixed cost of adjustment or some other friction can explain inertia and why change happens in jumps. However, fixed costs alone do not explain why it needs a crisis to implement change. Furthermore, in our model there is continuous (but slow) change in normal times followed by a jump in productivity when the crisis hits, which cannot be explained by fixed costs of adjustment. A jump of productivity in a crisis can be explained by a discontinuous decrease in workers' outside option in a crisis combined with a higher marginal utility of money at lower wages. This implies that workers have to be paid less to accept change in a crisis than in normal times. However, without a kink at the reference point (as in our model), the optimal adjustment in a crisis always uses both, a wage cut and an adjustment of work practices, so it does not explain downward rigidity of nominal wages. The downward rigidity of nominal wages can be explained by efficiency wages, insider-outsider dynamics, or contracting frictions, but these theories do not explain the described patterns of technological adaptation and change. Persistent productivity differences between firms can be explained by frictions in the dissemination of best practices, but this does not explain the other patterns described above. Thus, while our model is not the only explanation for the described phenomena, it is able to explain all of them simultaneously, based on the assumed loss aversion. Furthermore, it offers a new interpretation of expectation management that is consistent with rational expectations.

Our paper is related to three strands of the literature. First, there is an empirical literature on how competitive shocks affect productivity and productivity differences between firms. Bloom et al. (2014) report empirical evidence from the World Management Survey showing that there are large and persistent productivity differences across firms.[6] The authors show that higher total factor productivity is not only correlated with better management practices, but also with more intense competition. Performance increases in bad times, in particular, for low productivity firms. Bloom et al. (2019) show that the introduction of right-to-work laws (that weaken unions) in some states in the US is associated with improved managerial practices and efficiency increases. In Bloom et al. (2017) they also show that an increase in competition is associated with the introduction of better management techniques. Holmes and Schmitz

[6]See Syverson (2011) and Gibbons and Henderson (2013) for surveys of the literature on productivity differences between firms operating under similar conditions in narrowly defined industries.

(2010) survey case studies examining the behavior of firms that experienced dramatic changes in their competitive environment. They report that nearly all studies show that competitive shocks lead to increases in industry productivity. Plants that survive the competitive shock typically have large productivity gains. In addition, these gains often account for most of the overall industry gains. Using a rich data set on hundreds of markets in the ready-mix concrete industry in the US, Backus (2020) shows that an increase in competition in the ready-mix concrete industry has a positive causal effect on productivity (not driven by firm selection). All of this evidence is consistent with our model.

Second, our paper contributes to the small but growing literature on reference-dependent preferences in dynamic models. Eliaz and Spiegler (2014) develop a model of labor market dynamics in which workers have reference-dependent fairness preferences which gives rise to wage stickiness in recessions, similar to the downward rigidity of wages in our model. Pagel (2017) shows that the incorporation of dynamic loss aversion into a macro model can account for empirically observed consumption-savings patterns.[7] Pagel (2016) uses dynamic loss aversion to explain the observed equity premium volatility and the equity premium puzzle. Macera (2018) analyzes optimal incentive contracts in a dynamic moral hazard model with loss-averse agents. She shows that the principal backloads bonus compensation and pays a fixed wage in the present period if agents are sufficiently loss averse.[8] Alesina and Passarelli (2019) study the effects of loss aversion in electoral competition. Loss aversion gives rise to a status quo bias and an entrenchment effect: If there is an exogenous shock to voter preferences, the election outcome depends on the initial status quo. Furthermore, there are long-term cycles in policies with self-supporting movements to the right or to the left.[9] None of these papers

[7] Relatedly, Van Bilsen, Laeven, and Nijman (2020) find that with a backward-looking form of reference-dependence consumers delay painful consumption reductions.

[8] Relatedly, Herweg, Karle, and Müller (2018) examine the role of loss aversion on renegotiation in a buyer-seller setting. They emphasize the role of expectations, showing that if buyers do not expect renegotiation, then the parties may indeed not be able to renegotiate, even if the outcome is ex-post inefficient. Karle and Schumacher (2017) analyze the incentives of a monopolist to release ex-ante product information. They show that good information gives rise to an attachment effect if consumers are loss averse and adjust their expectations. In the context of auctions von Wangenheim (2021) and Balzer, Rosato, and von Wangenheim (2022) show that with dynamic reference-dependent preferences classic equivalences between dynamic and static auction formats break down. Rosato (2023) shows that a similar effect can explain empirically observed revenue declines in sequential auctions, since remaining bidders become less optimistic.

[9] Similarly, Lockwood and Rockey (2020) apply loss aversion to electoral competition in a representative democracy. They show that an incumbent reacts less strongly to a shift in voter preferences than challengers. Lockwood, Le, and Rockey (2022) study the interaction of loss aversion and incomplete recall in dynamic

applies dynamic reference dependence to organizational change. Furthermore, the main effect driving our results is new and does not play a role in the previous literature on dynamic loss aversion.

Finally, there is a large literature on downward nominal wage rigidity. Grigsby, Hurst, and Yildirmaz (2021) report for a large sample of 20 million workers in the US that only 2.5 percent of all workers who remain continuously employed with the same firm experience a nominal base wage cut in a given year, while about one third experienced no nominal wage adjustment. Fallick, Villar, and Wascher (2020) and Branten, Lamo, and Room (2018) show for the US and the EU, respectively, that downward nominal wage rigidity is highly prevalent, even during the Great Recession. Behavioral theories have emphasized the role of fairness (Fehr and Goette (2007), Fehr and Charness (2025)) and efficiency wages Akerlof (1982). There are other explanations of wage rigidity, including contractual frictions and insider-outsider theories, but none of them links it to organizational change.

The rest of the paper is organized as follows. The next section introduces the static version of the model. The model is "Coasean" in the sense that there are no informational or contractual frictions. The only friction is loss aversion. The model shows that loss aversion gives rise to inertia in normal times, but that parties will adjust toward material efficiency in a crisis. Section 3 sets up the dynamic infinite-horizon model, in which the reference point adjusts over time. The model shows that the principal offers a permanent wage increase to compensate for change in normal times, which gives rise to a quasi-rent that builds up over time. When a crisis hits, this quasi-rent is expropriated in order to implement drastic change. The model also shows how companies in the same industry using the same technology can have substantial productivity differences for extended periods of time. Section 4 introduces a principal-agent problem with a third party, the management, and shows how the principal can use expectation management to reduce the cost of implementing change. Section 5 concludes.

electoral competition.

2 A Coasean Model With Reference Dependence

There are two players, the owner of the firm (the principal, "she") and the workers, represented by a union (the agent, "he"), who negotiate on wages and the implementation of organizational change.[10] We focus on the effect of reference-dependent preferences of workers on the negotiation outcome, so we abstract away from any informational or contractual frictions. The parties can implement any change via efficient ("Coasean") bargaining.

We start out with a simple one-period model. There is a state of the world, $\theta \in \Theta \subset \mathbb{R}$, that represents the current state of technology. Workers have to take a costly action $x \in \mathbb{R}^+$ that adapts the organization to the state of the world. This gives rise to a gross profit $v(x, \theta)$ that accrues to the owner.

The principal's profit function is given by

$$\Pi = v(x, \theta) - w - K, \tag{1}$$

where $K \geq 0$ are all costs other than wages w. We assume that the gross profit function $v(x, \theta)$ is increasing and concave in x with $\frac{\partial v(x,\theta)}{\partial x} > 0$, $\frac{\partial^2 v(x,\theta)}{\partial x^2} < 0$, $\frac{\partial v(x,\theta)}{\partial \theta} > 0$, and $\frac{\partial v(x,\theta)}{\partial x \partial \theta} > 0$. Moreover, we assume $\lim_{x \to 0} \frac{\partial v(x,\theta)}{\partial x} = \infty$, $\lim_{x \to \infty} \frac{\partial v(x,\theta)}{\partial x} = 0$, and $\frac{\partial v(x,\theta)}{\partial x \partial \theta}$ strictly bounded away from zero in order to ensure interior solutions. The idea is that a higher state θ makes the change to more complex (i.e. higher) work practices x more profitable, but this change requires costly effort from workers. Without loss of generality we measure x by its cost, i.e., $c(x) = x$.

The utility function of the agent (the workers) is given by

$$U = w - x - \lambda[w^r - w]^+ - \lambda[x - x^r]^+, \tag{2}$$

where $[\cdot]^+ = \max\{\cdot, 0\}$. The agent's utility function is reference dependent. It consists of the material payoff, $w - x$, and the perceived experience of losses if the wage, w, is smaller than the reference wage, w^r, and/or if the action, x, is larger than its reference level, x^r. The parameter $\lambda > 0$ measures the degree of loss aversion.[11] The reference point (w^r, x^r) is a

[10]The agent could also be the manager of the firm. If the manager is loss averse and if organizational changes require adaptation costs, then the analysis and the results are exactly the same. We use "workers" or "union" not because we want to blame them for slow productivity growth but rather for concreteness as an example for the effects of loss aversion on change.

[11]There could be different parameters of loss aversion for wages w and action x. As long as they are positive, our qualitative results hold.

convex combination of the status quo, (w_0, x_0) and the rational expectation (w^e, x^e), i.e., the values of w and x that workers expect to be realized at the end of the period:

$$w^r = \alpha w_0 + (1-\alpha)w^e , \quad (3)$$

$$x^r = \alpha x_0 + (1-\alpha)x^e \quad (4)$$

where $\alpha \in [0,1]$ is the relative weight of the status quo.

For the baseline model we assume that organizational change is perfectly contractible. The principal makes a take-it-or-leave-it offer (w, x) to the agent (union/workers). If the agent rejects, the old contract (w_0, x_0) remains in place and the agent receives utility $U_0 \equiv w_0 - x_0 \geq 0$. Based on the offer and his rational expectation about his acceptance decision the agent forms his reference point.[12] We assume throughout that the agent expects to accept (w, x) if – given the expectation to accept – acceptance is optimal.[13] Then the agent decides whether to accept or reject the offer. After observing the agent's decision, the principal and thereafter the agent have the option to terminate the relationship. If one of them does, both parties get a utility of zero.

2.1 Inertia and Material Inefficiency in Normal Times

First, we consider the *normal case* where the old contract generates positive profits for the principal ($\Pi = v(x_0, \theta) - w_0 - K \geq 0$). This implies that the parties will continue their relationship even if the agent rejects the offer of the principal (which will be different in the case of a *crisis* to be discussed later). The agent correctly anticipates the negotiation outcome, so $(w^e, x^e) = (w, x)$ if he is going to accept the contract offer and $(w^e, x^e) = (w_0, x_0)$ otherwise. Because the agent anticipates that the relationship will be maintained even if he rejects the offer, his outside option utility in the normal case is given by $U_0 = w_0 - x_0 \geq 0$.

Thus, the principal's problem is:

[12]The idea that for $\alpha < 1$ the reference point (partly) adapts to the action chosen follows the logic of a choice-acclimating personal equilibrium in Köszegi and Rabin (2006).

[13]Without this assumption there are multiple equilibria, a common phenomenon in the literature on expectation-based reference points (Köszegi and Rabin (2006)). The other equilibria can be sustained by the workers' expectation to reject any wage offer below $\underline{w}$ with $x + \alpha\lambda(x - x_0) + U_0 < \underline{w} \leq x + \lambda(x - x_0) + U_0$. Hence, we assume that the firm can coordinate on the equilibrium in which the agent accepts change, if—given the reference point it induces—accepting change is optimal.

$$\max_{w,x}\{v(x,\theta) - w - K\} \tag{5}$$

subject to

$$U \;=\; w - x - \lambda[w^r - w]^+ - \lambda[x - x^r]^+ \;\geq\; U_0 \tag{6}$$

As a benchmark consider the case where there is no loss aversion ($\lambda = 0$), so workers are only concerned about their material payoff. In this case the principal offers $w = x + U_0$ and chooses the materially efficient level of x such that

$$\frac{\partial v(x^{ME},\theta)}{\partial x} = 1\ . \tag{7}$$

Note that x^{ME} is an increasing function of θ (because $v_{x\theta} > 0$). Moreover, it does not depend on the status quo (w_0, x_0) or the firm's cost K.

Consider now the case with loss aversion. In order to implement x the principal has to pay

$$\begin{aligned} w \;&=\; x + \lambda[\alpha w_0 + (1-\alpha)w - w]^+ + \lambda[x - \alpha x_0 - (1-\alpha)x] + U_0 \\ &=\; x + \alpha\lambda[w_0 - w]^+ + \alpha\lambda[x - x_0]^+ + U_0\ . \end{aligned} \tag{8}$$

We focus on the case where the principal wants to increase x (as compared to x_0) which implies that she also has to pay a higher wage to the workers. The case where the principal wants to decrease x is symmetric but less relevant, because θ, the state of technology can only go up. Thus, the principal maximizes

$$\Pi = v(x,\theta) - [x + \alpha\lambda(x - x_0) + U_0] - K\ .$$

The first order condition of this problem is

$$\frac{\partial v(x,\theta)}{\partial x} \leq 1 + \alpha\lambda\ . \tag{9}$$

with equality if $x > x_0$. Hence, the firm finds it optimal to increase x only if $\frac{\partial v(x_0,\theta)}{\partial x} > 1 + \alpha\lambda$.

The following proposition fully characterizes the optimal action x^* in normal times.

Proposition 1 (Optimal Contract, Normal Times)**.** *Suppose that the status quo contract* (w_0, x_0) *satisfies* $x_0 \leq x^{ME}(\theta)$ *and*

$$v(x_0,\theta) - w_0 - K > 0.$$

Define $\overline{x}(\theta)$ *implicitly by* $\frac{\partial v(\overline{x}(\theta),\theta)}{\partial x} = 1+\alpha\lambda$ *and note that* $\overline{x}(\theta) < x^{ME}(\theta)$*. The principal offers a contract* $(x^*(\theta), w^*(\theta))$ *to the agent that is given by*

$$x^*(\theta) \;=\; \begin{cases} x_0 & \text{if } x_0 \geq \overline{x}(\theta) \\ \overline{x}(\theta) & \text{if } x_0 < \overline{x}(\theta), \end{cases} \tag{10}$$

and

$$w^*(\theta) \;=\; \begin{cases} w_0 & \text{if } x_0 > \overline{x}(\theta) \\ w_0 + (1+\alpha\lambda)[\overline{x}(\theta) - x_0] & \text{if } x_0 < \overline{x}(\theta). \end{cases} \tag{11}$$

Proof. The proof follows directly from the arguments given in the text above. □

Proposition 1 shows that x^* differs from the materially efficient x^{ME}. Loss aversion drives a wedge between the marginal cost of the workers and the marginal benefit of the owner which induces the owner to stick to the status quo even if this is materially inefficient. Let $\overline{\theta}$ be implicitly defined by $\frac{\partial v(x_0,\overline{\theta})}{\partial x} = 1+\alpha\lambda$. For $\theta \leq \overline{\theta}$ there is full inertia.[14] But even if $\theta > \overline{\theta}$, $x^*(\theta)$ is strictly smaller than the materially efficient $x^{ME}(\theta)$. The next proposition shows how the distortion depends on the degree of loss aversion, reference point formation, and the status quo.

Proposition 2 (Comparative Statics)**.** *Suppose* $x_0 < x^{ME}(\theta)$*. The principal always implements less change than material efficiency requires, i.e.,* $x^*(\theta) - x_0 < x^{ME}(\theta) - x_0$*.*

Furthermore, we have

(a) If λ *(the degree of loss aversion) or* α *(the weight that workers put on the status quo in the formation of the reference point) increases, the amount of organizational change,* $x^* - x_0$*, decreases, i.e.*

$$\frac{\partial x^*}{\partial \lambda}, \frac{\partial x^*}{\partial \alpha} < 0 \quad \text{if } \theta \geq \overline{\theta} \tag{12}$$

$$\frac{\partial x^*}{\partial \lambda}, \frac{\partial x^*}{\partial \alpha} = 0 \quad \text{if } \theta < \overline{\theta}. \tag{13}$$

(b) An increase in λ *or in* α *widens the range of inertia, i.e.,* $\frac{\partial\overline{\theta}}{\partial\lambda}, \frac{\partial\overline{\theta}}{\partial\alpha} > 0$*, where* $\overline{\theta}$ *is implicitly defined by* $\frac{\partial v(x_0,\overline{\theta})}{\partial x} = 1+\alpha\lambda$*.*

[14] This implicitly assumes that $\theta \geq \underline{\theta}$ where $\underline{\theta}$ is defined by $x^{ME}(\underline{\theta}) = x_0$.

(c) An increase of x_0 increases $\overline{\theta}$.

Proof. See Appendix.

Without loss aversion ($\lambda = 0$) or with a reference point that is fully determined by rational expectations ($\alpha = 0$) the principal will implement the materially efficient outcome x^{ME}. With $\lambda, \alpha > 0$ there is inertia. The larger λ and α, the less change will be implemented and the larger the gap between the materially efficient action x^{ME} and the implemented action x^*. Furthermore, the larger the λ and α, the larger the range of inertia where the principal does not adjust x^* when θ increases. Finally, the range of inertia, i.e., $\overline{\theta}$, shifts upwards if the initial action x_0 increases.

2.2 A Parametric Example

Let $v(x, \theta) = \theta \ln x$. Then we have $\frac{\partial v(x,\theta)}{\partial x} = \frac{\theta}{x}$ and $\overline{\theta} = (1 + \alpha\lambda)x_0$. Thus we get $x^{ME}(\theta) = \theta$ and

$$x^*(\theta) = \begin{cases} x_0 & \text{if } \theta \leq (1 + \alpha\lambda)x_0 \\ \frac{\theta}{1+\alpha\lambda} & \text{if } \theta > (1 + \alpha\lambda)x_0. \end{cases} \tag{14}$$

Figure 1 illustrates the optimal choice of x^* by the principal for $\alpha = 0.5$, $\lambda = 1$, and $x_0 = 4$. These parameters imply $\overline{\theta} = (1 + \alpha\lambda)x_0 = 6$. The thick line depicts the optimal choice of $x^*(\theta)$ by the principal. If $\theta \leq \overline{\theta} = 6$, there is complete inertia. If $\theta > \overline{\theta}$, the principal adjusts x, but at a slope that is smaller than the slope of the materially efficient $x^{ME}(\theta)$. Thus, the larger the θ, the larger the gap between $x^{ME}(\theta)$ and $x^*(\theta)$.

2.3 The Effects of a Crisis

Suppose now that there is a crisis, i.e. a sudden shock to the firm's profits. The shock could be idiosyncratic (e.g. "Dieselgate" for Volkswagen) or it could affect the entire industry or economy (e.g. the Covid 19 pandemic). We model this exogenous decrease in profits as a shock to the firm's cost parameter K.[15] Consider a situation in which the cost parameter

[15] For tractability, we assume that the shock to profits does not affect the productivity of the worker.

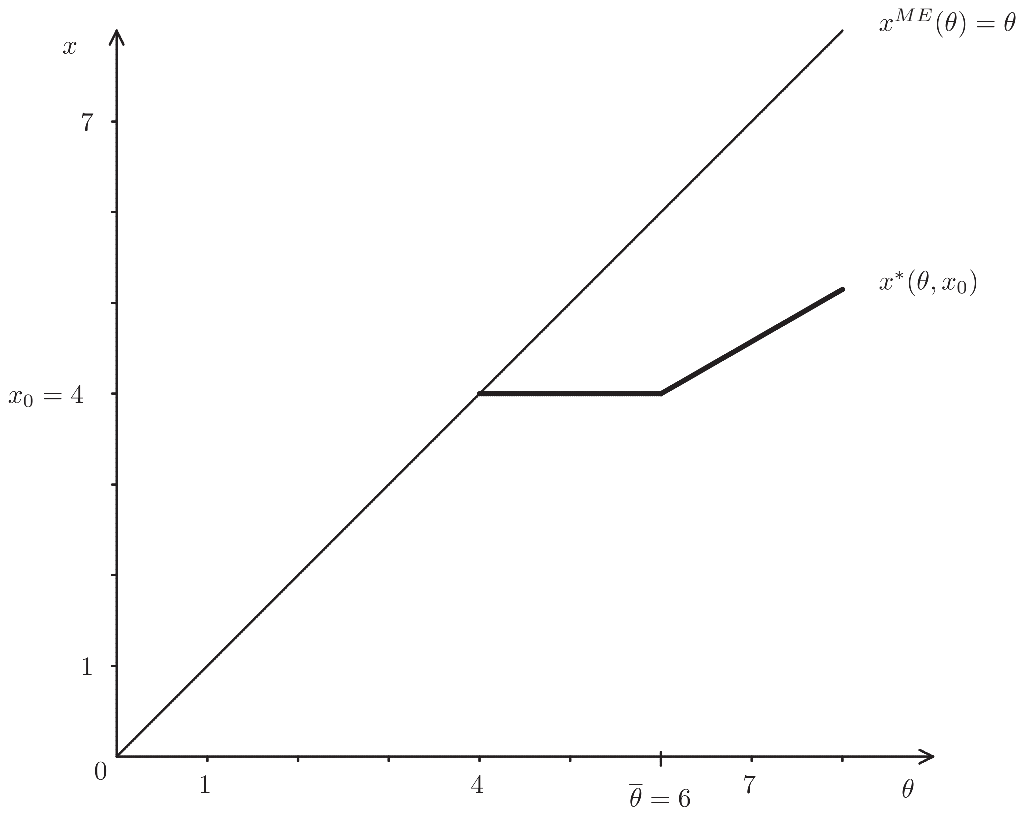

Figure 1: Organizational change as a function of θ with and without loss aversion.

K is such that the status quo contract (w_0, x_0) generates negative profits. Hence, the firm would rather terminate the relationship than continue with the old contract. This changes the workers' outside option which is now given by the utility level of unemployment that is normalized to 0. The next proposition shows how the firm uses the reduced outside option of the workers to change the terms of the contract.

Proposition 3 (Optimal Contract in a Crisis)**.** *Suppose that the status quo contract* (w_0, x_0) *satisfies* $x_0 \leq x^{ME}(\theta)$ *and*

$$v(x_0, \theta) - w_0 - K < 0.$$

Define $\hat{x}$ *implicitly by* $U(w_0, \hat{x}) = 0$ *and* $\overline{x}(\theta)$ *as in Proposition 1 by* $\frac{\partial v(\overline{x}(\theta), \theta)}{\partial x} = 1 + \alpha\lambda$.

1. *If* $\hat{x} \geq x^{ME}(\theta)$ *the firm offers a contract with* $x^* = x^{ME}(\theta)$.

2. *If* $\overline{x}(\theta) < \hat{x} < x^{ME}(\theta)$ *the firm offers a contract with* $x^* = \hat{x}$.

3. *If* $\hat{x} \leq \overline{x}(\theta)$ *the firm offers a contract with* $x^* = \overline{x}(\theta)$.

The offered wage w^* *always satisfies* $U(x^*, w^*) = 0$, *and the union accepts the offer.*

Proof. See the Appendix.

The workers know that they will lose their jobs if they reject the firm's offer. Therefore, any contract that offers them at least the utility level of unemployment will be accepted. The firm can use its improved bargaining position to either reduce wages w or to increase x. Thus, the firm pays for a higher x not with a higher wage (a gain), but with a smaller wage cut (a reduced loss). This changes the marginal rate of substitution between w and x. In a crisis, an increase in x comes at the same marginal cost to the worker as a wage cut of equal size. Because for $x < x^{ME}$ marginal productivity satisfies $\frac{\partial v(x_0,\theta)}{\partial x} > 1$. Hence, it is more efficient to increase x rather than decrease w. Additional wage cuts will only be implemented if the materially efficient level of x still generates positive utility for the workers at the status quo wage.

The third case in Proposition 3 considers the case where after increasing x to the workers' zero-utility threshold it is still smaller than the optimal level $\overline{x}(\theta)$ from Proposition 1. As in Proposition 1 the firm will then implement an additional increase of x up to the threshold $\overline{x}(\theta)$, which has to be compensated at a wage rate of $1 + \alpha\lambda$ to ensure the workers' consent.

Figure 2 illustrates Proposition 3. Consider a low status quo level of work practices, $x_0 < \overline{x}(\theta)$. Note that $\hat{x}$ linearly increases with the rent enjoyed by workers prior to the crisis. If this rent is small, the firm will increase x to the behaviorally efficient level $\bar{x}(\theta)$. If this rent is larger, the firm will increase x^* as much as possible while keeping wages constant, i.e. it sets $x^* = \hat{x}$. If the rent is so large that the firm could increase x even beyond the materially efficient level, it will stop at x^{ME} and reduce the workers' wages rather than increase x beyond x^{ME}.

In the static model the firm never wants to increase x beyond $x^{ME}(\theta)$. However, in the

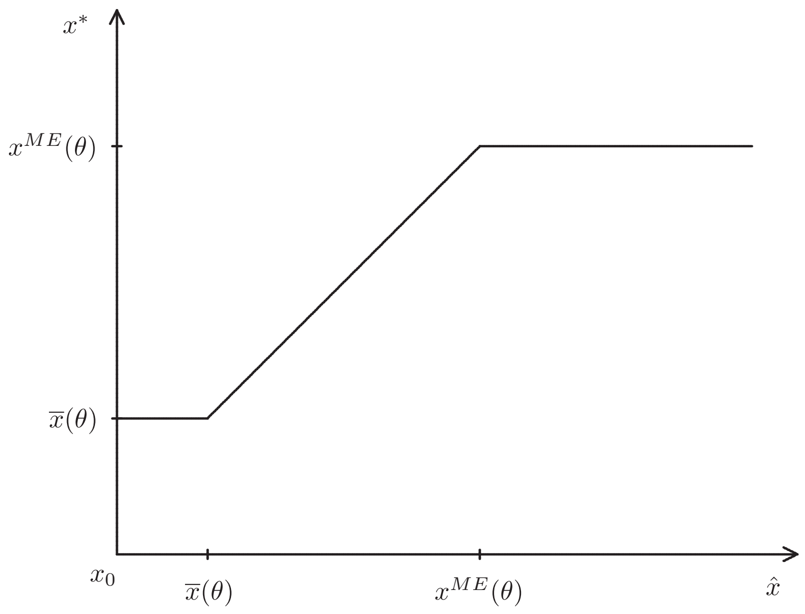

Figure 2: The Effect of a Crisis

dynamic model of the next section the firm might want to go beyond the materially efficient $x^{ME}(\theta)$ because it anticipates future increases of θ and wants to use the current crisis to prepare for these changes now already. However, it is difficult to adapt to a technological development that does not yet exist. Furthermore, this possibility gives rise to a number of uninteresting case distinctions. Therefore, the following assumption, made throughout the remainder of the paper, excludes this possibility:

Assumption 1. *Any contractible action must satisfy $x_t \leq x^{ME}(\theta_t)$, where $x^{ME}(\theta_t)$ denotes the efficient work practice for $\lambda = 0$.*

This assumption simplifies the analysis but does not affect the qualitative results of the following sections.

Proposition 3 shows that a crisis can induce a sudden jump in organizational efficiency. It also explains "why wages don't fall during a recession" (Bewley (1999)). In a recession, workers make concessions, but firms and workers will not negotiate wage cuts but will rather implement more change.

The idea that a crisis has a positive effect on economic efficiency goes back to at least Schumpeter (1942). This effect is partly due to a composition effect if less efficient firms are driven out of the market while more efficient firms stay in. But, as shown in many empirical studies (Holmes and Schmitz, 2010), it is also caused by a sudden increase in the efficiency of existing firm. We give a new perspective for this observed pattern: The potential for the efficiency increase existed already before the crisis, but the firm could not exploit it because of resistance to change. The effect of the crisis is to weaken this resistance of workers (and other stakeholders). This reduces the relative cost of implementing organizational change and makes change possible.

It is worth noting that while concessions increase material efficiency and may be necessary to sustain the firm, they are not "behaviorally" efficient if we incorporate the behavioral adaptation cost. Indeed, if $x_0 \in (\overline{x}, x^{ME}]$, a marginal increase of x decreases workers' utility at a rate of $1 + \alpha\lambda$, whereas firm profits only increase at a rate of $\frac{\partial v(x_0,\theta)}{\partial x} \in [1, 1 + \alpha\lambda]$. The reason is that utility is not perfectly transferable. Yet, transferring utility via adapting organizational practices is still more efficient than doing so via wage cuts.

3 Rents and Organizational Change over Time

The one-period model captures the main intuition as to why more change is implemented in a crisis than in normal times, but there are several open questions. First, the crisis has an effect only if workers enjoy a rent, i.e., $U_0 = w_0 - x_0 > 0$. Where does this rent come from? Second, it assumes that the reference point stays fixed after a change has been implemented. However, the reference point adjusts with some delay after the new contract becomes the new status quo. If the parties anticipate this, how does it change the optimal contract? Finally, rational parties anticipate that a crisis may occur. How do they prepare for it in the initial contract?

In this section we consider a dynamic model in which the reference point adjusts over time. We do not attempt to build a fully general model that allows for arbitrary cost shocks.

Instead, we focus on the most interesting case where the cost shock is such that the firm may survive only if the workers make concessions. We show that the qualitative insights of Section 2 carry over to the dynamic model. Furthermore, we show that workers accumulate over time a (quasi-)rent as a compensation for organizational change in the past. Parties rationally anticipate reference point adjustments and the possibility of a crisis and adjust contracts optimally, which strengthens our results. We delineate the long-term dynamics, illustrate how productivity differences may persist over time, and how a crisis may increase productivity and reduce the productivity gap between comparable firms in the same industry.

We proceed as follows: In Section 3.1 we set up the dynamic model. In Section 3.2 we start with the simpler case in which workers and the firm do not anticipate that a crisis may occur. We show how x is adjusted over time when there is no crisis and how the optimal contract is adjusted if an unanticipated crisis occurs. In Section 3.3 we show how productivity differences between firms facing the same technology may arise and how they are affected by an unanticipated crisis. Then, in Section 3.4, we show that our results persist and are even stronger with agents who form rational expectations. Finally, in Section 3.5 we discuss several extensions of the model.

3.1 The Dynamic Model

Time $t = 1, 2, ...$ is discrete with an infinite horizon. We start in $t = 1$ with some state θ_1 and some status quo contract (w_0, x_0) that satisfies $w_0 \geq 0$ and $x_0 \leq x^{ME}(\theta_1)$, where $x^{ME}(\theta_t)$ is the materially efficient action as defined in Assumption 1.

In every period a new state of the world (θ_t, K_t) materializes. The state of technology θ_t increases deterministically over time, i.e., $\theta_{t+1} > \theta_t$.[16] For simplicity we assume that the firm's cost realization $K_t \in \{0, K_t^h\}$ is equal to 0 in "normal" periods and that there is at most one "crisis" with a high cost shock $K_t^h > 0$. Conditional on zero costs in all previous periods, there is a crisis in the next period with probability $\mu > 0$. We are interested in the case where K_t^h is sufficiently large that the firm prefers to terminate the relationship if workers do not agree to

[16] We assume implicitly that the growth in θ is bounded such that the present value of revenues $\sum_{t=0}^{\infty} \delta^t v(x_t, \theta_t)$ remains finite for all work practices, such that present values are well defined.

make concessions (Assumption 2 below). A more general structure of cost shocks would give rise to many case distinctions which do not yield interesting additional insights.

Adapting to new work practices, again, requires costly effort from workers in the period when the change is implemented.[17] Without loss of generality, we measure the magnitude of a change in x by its cost.

The timing in each period is as follows. First, both parties observe the new state of the world. Then, the firm makes a take-it-or-leave-it offer (w_t, x_t) to the workers. The workers may accept or reject the offer. If they reject, the current contract (w_{t-1}, x_{t-1}) remains in place. After observing the workers' decision, both parties may or may not terminate the relationship.

Each party maximizes their expected utility, where future utility is discounted at a common discount factor $\delta < 1$. Hence, for an observed state (θ_t, K_t) and a contract (w_{t-1}, x_{t-1}) inherited from the previous period the firm's expected future profit evaluated at the beginning of period t is given by

$$\Pi((w_{t-1}, x_{t-1}), (\theta_t, K_t)) = \sum_{s=t}^{\infty} \delta^{s-t} \cdot \mathbb{E}_t[v(x_s, \theta_s) - w_s - K_s]. \tag{15}$$

The workers' reference point in period t is a convex combination of the status quo contract (w_{t-1}, x_{t-1}) and the correctly anticipated negotiation outcome (w_t, x_t). Again, workers receive reference-dependent utility by comparing the negotiation outcome in period t to the reference point in that period. Note that in the dynamic model there is a difference between a wage cut and the cost to adapt to a new work practice. The adaptation cost is incurred in one period only, while the wage cut has a permanent effect on all future wages. Therefore, we assume that the loss experienced with a wage cut is proportional to the present value of the permanent wage reduction.[18]

[17]In an earlier version of the paper we considered a model in which changes to "higher" work practices require permanently higher effort. In this case, the wage payment in each period has to compensate for this additional cost. For such a model, all results and intuitions remain qualitatively unchanged.

[18]Note that while the wage cut applies in every period, the experience of the loss is felt only in the period in which the wage cut is implemented. Thereafter, the reference point adjusts to the new wage level.

Thus, the workers' continuation utility is

$$
\begin{aligned}
U_t(w_t, x_t | w_{t-1}, x_{t-1}) &= \sum_{s=t}^{\infty} \delta^{s-t} \cdot \mathbb{E}_t \left[w_s - |x_s - x_{s-1}| - \frac{\lambda}{1-\delta}[w_s^r - w_s]^+ - \lambda |x_s - x_s^r| \right] \\
&= \sum_{s=t}^{\infty} \delta^{s-t} \cdot \mathbb{E}_t \left[w_s - |x_s - x_{s-1}| - \frac{\lambda\alpha}{1-\delta}[w_{s-1} - w_s]^+ - \lambda\alpha |x_s - x_{s-1}| \right].
\end{aligned} \tag{16}
$$

We restrict attention to Markov perfect equilibria in the sense that any offer (w_t, x_t) and each party's subgame-perfect decision to end or continue the relationship depend only on the current state, i.e. (θ_t, K_t) and the current reference point. In particular, both parties cannot commit to any future path of actions or transfers. Moreover, we assume that if one party is indifferent in its decision, it breaks ties in favor of the other party.

3.2 An Unanticipated Crisis

In this section, we analyze the more tractable case in which the players anticipate that the reference point will change, but they do not anticipate that a crisis will occur with positive probability. The more general case in which the crisis is rationally anticipated is covered in Section 3.4.

Suppose that the cost realization will be zero in each period, that is, $\mu = 0$. Ignoring participation constraints the (behaviorally) efficient contracts solve:

$$
\max_{(w_t, x_t)_{t \geq 1}} W = \sum_{t=1}^{\infty} \delta^{t-1} \big(v(x_t, \theta_t) - |x_t - x_{t-1}| - \frac{\lambda\alpha}{1-\delta}[w_{t-1} - w_t]^+ - \lambda\alpha |x_t - x_{t-1}| \big).
$$

Note that the wage affects efficiency only in the case of wage cuts. Falling wages create an inefficiency due to the behavioral cost. Hence, any behaviorally efficient choice of work practices that is accompanied by a weakly increasing wage schedule is a behaviorally efficient solution. The following lemma characterizes the behaviorally efficient sequence of work practices in the dynamic model:

Lemma 1. *Define $\overline{x}(\theta)$ implicitly by the level of work practice that satisfies $\frac{\partial v(x,\theta)}{\partial x} = (1 - \delta)(1 + \alpha\lambda)$. For $\mu = 0$ the work practices $(x_t^*)_{t \geq 1}$ of any behaviorally efficient contract satisfy*

$x_{t+1}^* \geq x_t^*$ *with*

$$x_t^* = \begin{cases} x_{t-1} & \text{if } \; x_{t-1} \geq \overline{x}(\theta_t) \\ \overline{x}(\theta_t) & \text{if } \; x_{t-1} < \overline{x}(\theta_t). \end{cases}$$

Proof. See the Appendix.

Lemma 1 illustrates the trade-off between implementing change immediately and delaying it by one period. If a change of one unit is delayed by one period, the firm looses $\frac{\partial v(x,\theta)}{\partial x}$ in the current period. At the same time, it delays the adaption cost by one period, so it saves $(1 + \alpha\lambda)$ this period, but it has to incur it in the next period, where the cost is discounted with δ. Thus, the saving is $(1 - \delta)(1 + \alpha\lambda)$.

It is insightful to compare the behaviorally efficient work practices with the materially efficient level when $\lambda = 0$. If $\lambda = 0$ and work practices x_0 start at a sufficiently low level, then the optimal schedule of work practice satisfies

$$\frac{\partial v(x^{ME}, \theta)}{\partial x} = 1 - \delta.$$

Because $\overline{x}(\theta_t) < x^{ME}(\theta_t)$, the path of behaviorally efficient work practices is reminiscent of that chosen by the firm in Proposition 1. It is weakly increasing, because θ_t is going up each period. If we start at the materially efficient level, then there is, again, a region of inertia. Only if the mismatch between the state θ_t and the current work practice x_{t-1} is sufficiently large, it is optimal to change the work practice, but it is always smaller than the materially efficient level.

We are now ready to characterize the equilibrium in the benchmark case where a crisis is not anticipated ($\mu = 0$). The firm will simply implement the behaviorally efficient schedule of work practices and spread the necessary wage increase over all future periods.

Proposition 4. *If $\mu = 0$ the following is the unique equilibrium.*

1. *In period $t \geq 1$ the firm offers the contract (w_t^*, x_t^*), where x^* is the behaviorally efficient choice of work practices characterized in Lemma 1 and $w_t^* = w_{t-1} + (1-\delta)(1+\alpha\lambda)(x_t^* - x_{t-1})$.*

2. *The workers accept.*

To see the intuition for this result, recall that the behaviorally efficient x is weakly increasing over time. In period t the contract (w_{t-1}, x_{t-1}) constitutes the workers' outside option, so the firm has to compensate the workers for any change in work practice in order to guarantee the same utility as under contract (w_{t-1}, x_{t-1}). The necessary additional compensation of $(1+\alpha\lambda)(x_t^* - x_{t-1})$ covers both, the material and behavioral cost of adaptation. Note that it would be inefficient to fully compensate the worker in period t for the arising cost. Such a compensation would lead to a decline in wages in $t+1$ after the reference level of wages has adjusted upwards, causing the experience of a loss by workers that needs to be compensated. To avoid feelings of losses associated with future wage reductions any behaviorally efficient compensation scheme must involve weakly increasing wages.[19] The central observation is that the compensation must be spread out evenly over all future periods. Indeed, the present value of a permanent payment $(1-\delta)(1+\alpha\lambda)(x_t^* - x_{t-1})$ is

$$\sum_{s=t}^{\infty} \delta^{s-t}(1-\delta)(1+\alpha\lambda)(x_t^* - x_{t-1}) = (1+\alpha\lambda)(x_t^* - x_{t-1}).$$

In our model there is no other feasible compensation schedule to implement the behaviorally efficient schedule of work practices. Since the firm has no commitment power it cannot backload the compensation to the future. As argued above, frontloading the compensation is inefficient, as this would imply that wages have to fall in some future periods. Thus, implementing the efficient work practices in this way is the best the firm can do, because it generates the highest possible joint surplus, but leaves the workers with only the utility of their outside option.

Proof. Follows directly from the text. □

Note that the optimal contract in period t results in a utility loss for workers in period t, but a permanent utility increase in all future periods. Thus, starting in period $t+1$ workers enjoy a quasi-rent. Because the compensation is linear in the increase of work practices,

[19] An immediate full compensation is equally optimal only if the reference point does not adjust at all. For example, if the change in x is implemented very quickly, and if the firm compensates workers with a one-time bonus that has no effect on their reference point and does not change the status quo contract, then a one-time bonus would also be optimal. However, if the reference point adjust only a little bit, a permanent wage increase is strictly better.

the workers' wage and their quasi-rent at each point in time can be easily derived from the compensations for the sum of past adaptations $\sum_{s=1}^{t-1}(x_s - x_{s-1}) = (x_{t-1} - x_0)$:

$$w_{t-1} = w_0 + (1-\delta)(1+\alpha\lambda)(x_{t-1} - x_0), \tag{17}$$

and

$$\begin{aligned} U_t^*(w_t, x_t) = U_t(w_{t-1}, x_{t-1}) &= \sum_{s=t}^{\infty} \delta^{s-t}\big[w_0 + (1-\delta)(1+\alpha\lambda)(x_{t-1} - x_0)\big] \\ &= U_1(w_0, x_0) + (1+\alpha\lambda)(x_{t-1} - x_0), \end{aligned} \tag{18}$$

where $U_1(w_0, x_0) = \sum_{s=0}^{\infty} \delta^s w_0 = \frac{w_0}{1-\delta}$.

The firm's equilibrium profit in period t consists of the discounted sum of future surpluses minus the utility left to the workers, i.e.,

$$\Pi_t^* = \sum_{s=t}^{\infty} \delta^{s-t}\big(v(x_s^*, \theta_s) - (1+\alpha\lambda)(x_s^* - x_{s-1}^*)\big) - (1+\alpha\lambda)(x_{t-1}^* - x_0) - U_1(w_0, x_0). \tag{19}$$

Suppose now that in some period t there is a cost shock that decreases the present value of the firm's expected future profits by K_t^h. If the firm's value remains positive then the shock has no impact on the firm's optimization problem. The interesting case is when the cost shock induces a negative firm value for the status quo contract. The following assumption restricts attention to this case and is maintained throughout the remainder of the paper.

Assumption 2. *In every period the potential cost shock satisfies $K_t^h > \Pi_t^*$, where Π_t^* is defined in(19).*

Assumption 2 implies that if there is a cost shock then the value of the firm will become negative if workers receive their status quo utility, even if the most efficient contract is implemented.[20] Hence, workers are willing to make concessions to prevent the firm from closing down. The unemployment utility of zero constitutes the new threat point in the negotiations. Note that we allow for the possibility of the cost shock being so high that the firm cannot be rescued even if workers give up their quasi-rent. In this case the firm closes down and workers become unemployed.

[20]Thus, a fortiori, the value of the firm becomes negative for any other feasible contract as well.

The following proposition is a straightforward adaptation of Proposition 3 to the dynamic case.

Proposition 5. *Suppose that Assumptions 1 and 2 hold and there is a crisis in period t. Define $\hat{x}$ implicitly by $U_t(w_{t-1}, \hat{x}) = 0$, and $\overline{x}(\theta)$ as in Lemma 1 by $\frac{\partial v(\overline{x}(\theta),\theta)}{\partial x} = (1-\delta)(1+\alpha\lambda)$.*

1. *If $\hat{x} \geq \overline{x}(\theta)$ the firm offers a contract with $x^* = \min\{\hat{x}, x^{ME}(\theta)\}$.*
2. *If $\hat{x} < \overline{x}(\theta)$ the firm offers a contract with $x^* = \overline{x}(\theta)$.*

The offered wage w^* always satisfies $U(x^*, w^*) = 0$, and the union accepts the offer. The firm decides to continue the relationship if and only if its expected profit from the above contract is nonnegative.

Proof. See the Appendix.

Again, the firm finds it more profitable to increase work practices rather than decrease wages. Decreasing the wage or increasing the work practice by one unit has identical effects on the workers' expected utility. For the firm, increasing the work practice is even more appealing than in the static case: since θ is growing over time, a higher level of work practices today avoids costly adaptations in the future.

3.3 Long-Term Dynamics and Persistent Productivity Differences

We now illustrate the long-term dynamics and show how a crisis can help close productivity gaps between firms. We continue with our previous example of $v(x, \theta) = \theta \ln(x)$. The following Figure 3 depicts the materially efficient line $x^{ME}(\theta) = \frac{\theta}{1-\delta}$ as a solid line and the behaviorally efficient line $\overline{x}(\theta) = \frac{\theta}{(1-\delta)(1+\alpha\lambda)}$ as a dotted line. We consider two firms that have access to the same production technology. The productivity of the technology is summarized by θ, which grows over time. Firms are founded at different points in time. Firm 1 was founded when $\theta = 4$ and Firm 2 when $\theta = 7.5$, respectively. Thus, Firm 1 starts production at the materially efficient point A. The thick solid line depicts the transition path as θ_t grows over time. Firm 1 is in the area of inertia until point B is reached. From point B onward, the contract follows the behaviorally efficient (dotted) line. The firm implements higher work

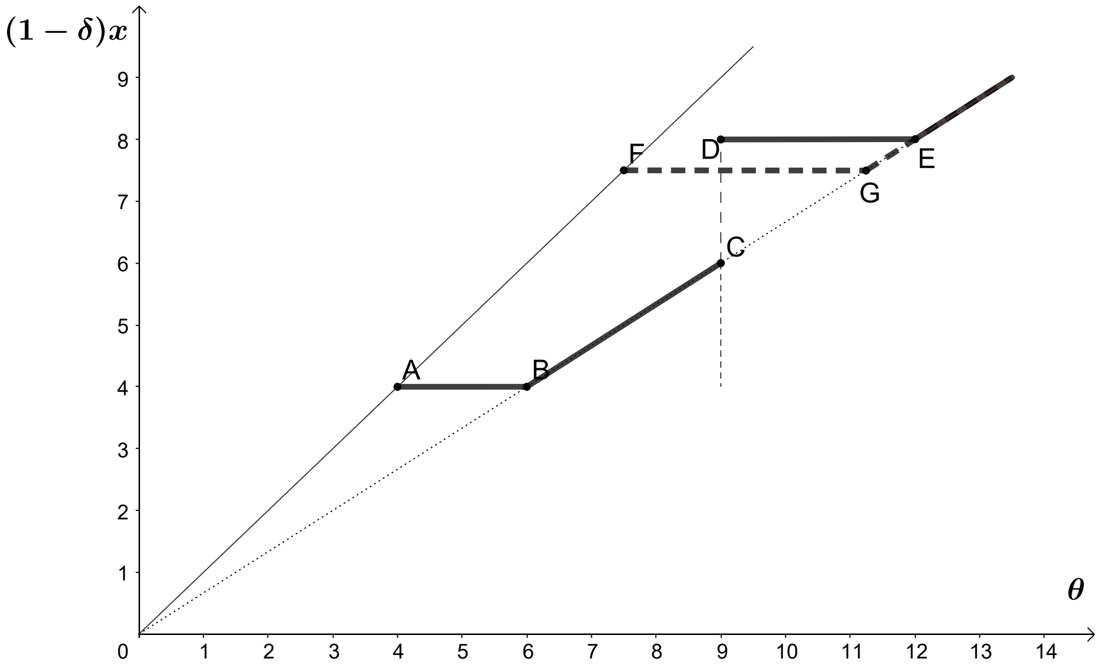

Figure 3: $v(x, \theta) = \theta \ln x$, $\alpha\lambda = 0.5$, firms created at $\theta = 4$ and $\theta = 7.5$, crisis at $\theta = 9$

practices each period, but these increases are smaller than the materially efficient increases. Because the compensation for the adaptation cost is stretched out over time, workers build up a quasi-rent.

At $\theta = 7.5$ firm 2 is founded. Firm 2 also starts production at a materially efficient point (F) and follows the transition path along the thick dashed line. Note that there is a substantial productivity gap between the two firms. This gap gradually closes as Firm 1 slowly implements higher work practices, but it may persist for a long time.

When $\theta = 9$ is reached, the economy is hit by a crisis that threatens the profitability of Firm 1, but not of the more efficient Firm 2. By then Firm 1 is at point C, and the workers have accumulated a rent of $(1 + \alpha\lambda)$ times the magnitude of total work practice increases, illustrated by the length of the dashed black line. Firm 1 will use the threat of unemployment to expropriate this quasi-rent by increasing work practices from point C to point D. If Firm 1 is profitable at point D, it will stay in business and continue producing, now again in the area of inertia until it hits point E.

The more efficient Firm 2 is not threatened by bankruptcy, even though it is hit by the same cost shock. Hence, the crisis leaves the implemented work practices unaffected. Hence, the productivity gap between the two firms closes discontinuously in a crisis, since Firm 1 is able to use the crisis to implement more efficient work practices, while Firm 2 is not. If the respective quasi-rents of Firm 1 are sufficiently high, as in the displayed example, Firm 1 will implement even more modern work practices than Firm 2.

The figure assumes equal levels of loss aversion of workers across firms. This does not necessarily have to be the case. Gächter, Johnson, and Herrmann (2022) find that loss aversion tends to increase in age, income, and wealth. This suggests that if younger firms employ a younger workforce, their workers suffer to a lesser degree from loss aversion than their older colleagues at older firms. In this case the inertia region of old firms is larger than that of young firms since workers from older firms are more reluctant to agree to organizational change. This could be another source of persistent productivity differences across firms.

3.4 The Equilibrium with Rational Expectations

Now we analyze the case in which both players correctly anticipate that the crisis may occur with probability μ each period. We show that all insights continue to hold and the range of inertia even widens.

We begin the analysis by noting that Assumption 2 implies that the parties cannot affect the probability of the crisis. Hence, in the contracting game, both parties treat the probability of a crisis as exogenous.

Second, notice that the optimal reaction to a crisis for a given status quo contract is fully analyzed in Proposition 5. Indeed, due to the restriction to Markov strategies, the contract offer in a crisis depends only on the current reference point and the states θ_t and K_t^h. Hence, the problem reduces to finding the principal's optimal contract offer in normal periods, given that both parties correctly anticipate the effects of their contract on the adaptation in a potential crisis.

We start by calculating the necessary wage compensation to implement an increase from x_{t-1} to x_t in a normal period. The firm has to compensate the worker for the adaptation cost of $(1+\alpha\lambda)(x_t-x_{t-1})$, which consists of both, material and behavioral adaptation cost. Again, the compensation must be spread equally over the current and all future normal periods. However, as compared to the case with an unanticipated crisis, both parties now anticipate that the crisis may occur, in which case workers will lose their quasi-rent. Hence, both parties understand that the de facto discount factor for the stream of per-period compensation now is $\delta(1-\mu)$. With this adjusted discount factor a constant per-period compensation of $(1-\delta(1-\mu))(1+\alpha\lambda)(x_t - x_{t-1})$ leads to a present value utility of

$$\sum_{s=t}^{\infty}(\delta(1-\mu))^{s-t}(1-\delta(1-\mu))(1+\alpha\lambda)(x_t^* - x_{t-1}) = (1+\alpha\lambda)(x_t^* - x_{t-1}).$$

Hence, wages in normal times satisfy

$$w_t = w_{t-1} + (1-\delta(1-\mu))(1+\alpha\lambda)(x_t - x_{t-1}).$$

We modify Equations (17) and (18) accordingly. Iterating the wage formula leads to

$$w_t = w_0 + (1-\delta(1-\mu))(1+\alpha\lambda)(x_t - x_0), \tag{20}$$

and

$$U_t(w_{t-1}, x_{t-1}) = \frac{w_0}{1-\delta(1-\mu)} + (1+\alpha\lambda)(x_{t-1} - x_0),$$

where the first summand corresponds to $U_1(w_0, x_0)$ and the second summand corresponds to the present value of the quasi rent.

Again, over time, as no crisis occurs, the workers build up a quasi-rent that stems from the compensation for past adaptations of work practices. Note that since workers anticipate the loss of their quasi-rent in a crisis, the per-period compensation they demand increases in the probability μ that a crisis will occur.

As established in Proposition 5, in a crisis the firm will take away the workers' utility to implement higher work practices. The following Lemma shows by how much the firm can maximally increase x in a crisis when keeping the wage constant.

Lemma 2. *Consider a status-quo contract (w_{t-1}, x_{t-1}) for which w_{t-1} satisfies (20). The increase of work practice Δx_t that satisfies $U_t(w_{t-1}, x_{t-1} + \Delta x_t) = 0$ is given by*

$$\Delta x_t = \frac{w_0}{(1-\delta)(1+\alpha\lambda)} + \frac{1-(1-\mu)\delta}{1-\delta}(x_{t-1} - x_0). \tag{21}$$

Proof. See the Appendix. □

The next proposition characterizes the optimal schedule of work practices that the firm will implement in normal times in anticipation of how the contract will change if a crisis hits.

Proposition 6. *Suppose there is no crisis until period t. There is a threshold $\tilde{x}(\theta_t) \leq \overline{x}(\theta_t)$ such that the optimal work practice x_t implemented in period t is given by*

$$x_t = \begin{cases} x_{t-1} & \text{if } x_{t-1} > \tilde{x}(\theta_t) \\ \tilde{x}(\theta_t) & \text{if } x_{t-1} \leq \tilde{x}(\theta_t). \end{cases}$$

Proof. See the Appendix. □

Comparing Proposition 6 to Proposition 4 shows that if parties rationally anticipate that a crisis may hit with probability $\mu > 0$, there will be even more inertia than if they do not anticipate a crisis, i.e. $\tilde{x}(\theta_t) \leq \overline{x}(\theta_t)$. Because the firm anticipates that it will become cheaper to adjust work practices in a crisis, it delays its adaptation.

The proof of Proposition 6 requires a few case distinctions and is somewhat involved, but the intuition can be illustrated graphically with the logarithmic example that we used in Section 3.3. Figure 4 reconsiders the example of Figure 3 with $\mu = 0.4$. The thin solid line depicts the boundary of the inertia region with players who form rational expectations. Note that it is further to the right than the behaviorally-efficient thin dashed line.

To see that the anticipation of the crisis increases inertia, two cases have to be distinguished. First, if the cost shock K_t^h is sufficiently large, the workers' accumulated rent will not suffice to make the firm profitable and the principal will terminate the relationship. The anticipation of bankruptcy makes the implementation of higher levels of work practices less appealing, as the firm will not benefit from this adaptation after bankruptcy.

Second, and more interestingly, the inertia area also widens when workers' concessions suffice to keep the firm in business. There is more inertia than without anticipation of the crisis because the principal knows that work practices will be adapted in a crisis anyway and therefore delays adaptation in normal times. Intuitively, this behavior helps the principal to keep the level of work practices closer to the dashed behaviorally efficient line $\overline{x}(\theta)$. It optimally

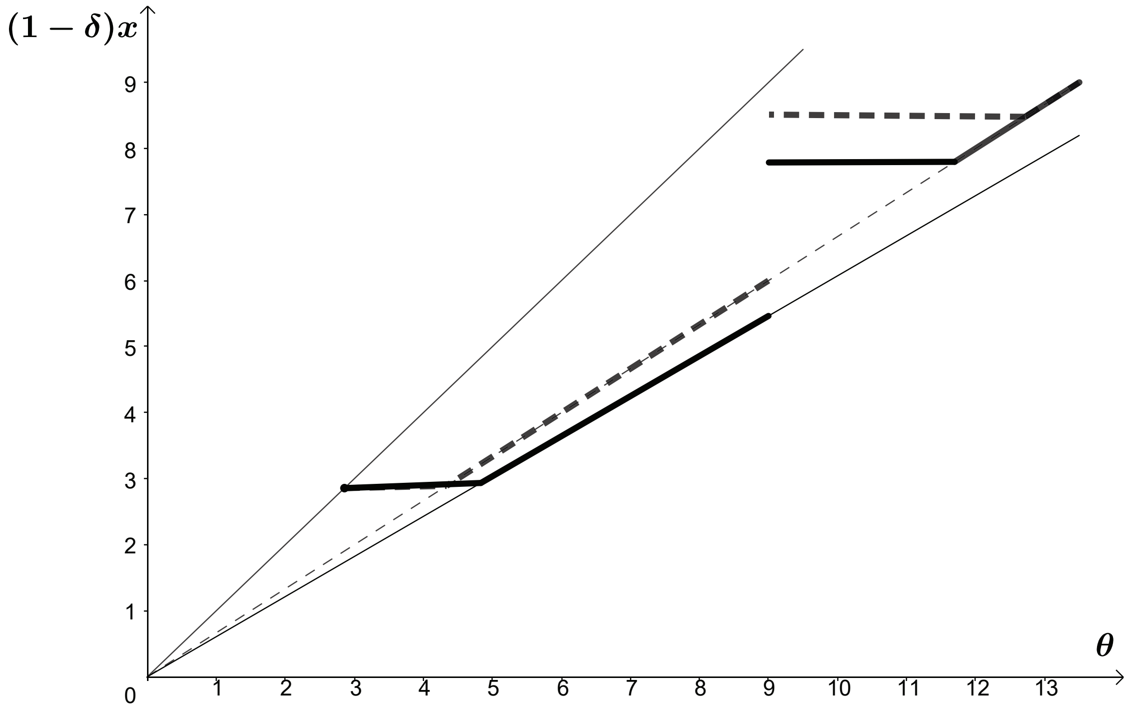

Figure 4: Logarithmic Example with Forward-Looking Players

solves a trade-off between "too low" work practices in normal periods and "too high" work practices in a crisis as compared to the behaviorally efficient level of work practices. The thick solid line depicts the optimal transition path of work practices if a potential crisis is anticipated, as compared to the case of an unanticipated crisis (dashed solid line).

3.5 Extensions

The qualitative results of the model are robust with respect to several extensions and variations of modeling choices.

Bargaining power. The model assumes, for simplicity, that the firm has all the bargaining power and can make a take-it-or-leave-it offer to workers. However, all qualitative results hold for any distribution of bargaining power. To see this, suppose that the firm wants to propose a change of x in a normal period (no crisis). Because utility is perfectly transferable from the firm to the workers through an increase in wages, the parties will agree to the behaviorally efficient work practice $x^*(\theta)$, which is independent of the distribution of bargaining power. However, the stronger the bargaining position of workers, the higher the permanent wage increase and the larger the quasi-rents they get.[21]

Thus, if a crisis threatens the viability of the firm, there are more rents that can be expropriated to implement change. However, more bargaining power for workers means that they also get some of the surplus in the crisis. For example, if workers have all the bargaining power, they agree to organizational change only up to the point where the firm just survives. This limits the rents that can be used to implement change.

Inflation. How does inflation affect the results of the model? The reference wage is a combination of the nominal status quo wage and the rationally anticipated nominal wage for the current period. If the firm offers a permanent nominal wage increase, workers anticipate that this will be eroded by inflation. Thus, as a compensation for the temporary adaptation cost they require a larger permanent wage increase. Hence, for any organizational change workers are compensated with a higher wage increase in the beginning that depreciates in real terms with the inflation rate over time. Note, however, that there are new adaptations in every period as θ_t increases over time, giving rise to additional quasi-rents. How the sum of these rents evolves over time depends on the time path of θ_t, but this does not affect the qualitative results of the model.

Permanently higher effort costs The model assumes that the material cost of implementing an organizational change is incurred in a single period. However, in some applications change involves permanently higher effort of workers. We analyzed this case in an earlier version of the paper, Schmidt and von Wangenheim (2023). There the material cost of a higher action is permanent, whereas the psychological cost associated with changing to a higher action or a lower wage occurs only in the period in which the change is implemented. Concretely, for a status quo contract (w_{t-1}, x_{t-1}) the workers' utility in period t from a contract (w_t, x_t)

[21] The size of the permanent wage increase depends also on whether the crisis is anticipated. Similar as in 3.4 workers understand that their compensation for past adaptations will be expropriated in a crisis. Hence, if the crisis is anticipated, workers demand a higher per-period wage increase to compensate for this.

is then given by

$$w_t - x_t - \lambda\alpha[w_{t-1} - w_t]^+ - \lambda\alpha[x_t - x_{t-1}]^+$$

Whenever the firm wants to increase the workers' action, workers must be permanently compensated for the higher material cost of effort. The additional compensation for the temporary psychological adaptation cost of $\lambda\alpha(x_t - x_{t-1})$ is again spread out over time in order to avoid wage cuts, i.e.

$$w_t = w_{t-1} + (x_t - x_{t-1}) + (1-\delta)\lambda\alpha(x_t - x_{t-1}).$$

Notice that the marginal cost of adaptation for the firm is $1 + (1-\delta)\lambda\alpha$. Hence, in normal times there is inertia whenever $\frac{\partial v(x,\theta)}{\partial x} < 1 + (1-\delta)\lambda\alpha$. If change is implemented it remains below the materially efficient level, and the compensation for past adaptation gives rise to a quasi-rent that can be appropriated in a crisis. Hence, all our qualitative results continue to hold for permanent changes in effort costs.

4 Expectation Management

In the preceding sections the implementation of organizational change is deterministic. However, in the real world it is often uncertain whether change is going to succeed, and managers have to work hard to push change through. Management textbooks and practitioners emphasize that expectation management is of crucial importance.[22] If a firm wants to implement change, it must convince all involved parties that change is inevitable and that everybody should prepare for it rather than resist it. Note that this is particularly difficult in normal times. In a crisis everybody knows that the firm will go bankrupt if change is unsuccessful, so everybody has a strong incentive to avoid this. The more interesting case is implementing change when change increases profits but the firm will also survive without it. In this case, it is more difficult to convince workers that change is inevitable.

In this section, we model the implementation of change explicitely and show how expectation management can work in normal times. As in the previous sections, parties form rational expectations. In the extended model, workers' consent to change is still necessary

[22] See e.g., Kotter (1995) and Burke (2017).

but no longer sufficient for implementing change. There is a manager whose job is to push change through. The harder the manager works, the more likely the change will succeed. Effort is costly to the manager.[23] The (moral hazard) problem is that the manager's effort is not observable to the owner of the firm and, therefore, cannot be contracted upon. Thus, the owner of the firm must offer an incentive contract to the manager. We show that this incentive contract serves as a commitment device that can be used to manage workers' expectations, lowering their resistance to change.

To keep the model tractable, we consider only the static model and take the amount of organizational change, Δx, as exogenously given. If there is a change of $\Delta x > 0$, the owner's gross profit increases by $\Delta v > \Delta x$. Furthermore, we restrict attention to the case where the firm survives even if change is not successful (no crisis).[24] Change has to be implemented by a manager who has to spend effort to increase the probability that organizational change will be successful. If he is unsuccessful, there is no change. We assume that the manager chooses the probability of success, p, at effort cost $c(p)$. In order to derive a closed-form solution, we consider a quadratic cost function $c(p) = \frac{c}{2}p^2$.

The owner of the firm, the principal, has to incentivize the manager to spend effort. The manager's chosen probability of success is unobservable to the principal and cannot be contracted upon. Change, however, is contractible, so the principal can offer a bonus payment b if the manager is successful in implementing it. The manager's outside option utility is normalized to 0 and he is wealth constrained, so he has to get a wage that is greater than or equal to 0 in both states of the world. We assume again that only workers suffer from loss aversion (as in Section 2) and that all other parties are loss and risk neutral and maximize the expected value of their payoff.

Note that in Section 2 everything is deterministic, while in this section the outcome is stochastic. For this stochastic environment the workers' reference point is a convex combination of the status quo and the ex ante expected value of the outcome.[25] Hence, if change Δx

[23]Formally, our model in Section 2 is a special case of the model in this section where the manager's effort costs are zero such that she always pushes change through with probability one.

[24]There is no problem in endogenizing the size of Δx and to consider the case of a crisis, but it complicates the exposition significantly without adding any new insights.

[25]Köszegi and Rabin (2006) assume that the decision maker compares the realized outcome to each possible outcome in the ex-ante distribution rather than the expected value. In our simple model their specification

is implemented with probability p the workers' reference point in the action dimension is[26]

$$x^r = \alpha x_0 + (1-\alpha)\big((1-p)x_0 + p(x_0+\Delta x)\big) = \alpha x_0 + (1-\alpha)[x_0 + p\Delta x].$$

The time structure of the model is as follows: At stage 1 the principal makes a take-it-or-leave-it contract offer to the manager. Workers observe this contract and form rational expectations. At stage 2 the principal offers workers a wage payment w to accept the change if the manager is successful.[27] At stage 3 the manager chooses the probability of success. At stage 4 nature determines whether change is successful and payments are made.

The principal's problem is a standard moral hazard problem with a risk-neutral but wealth-constrained manager. It is straightforward to show that the owner will offer the manager a wage of 0 if he fails and a bonus $b \geq 0$ if he is successful.

Suppose that the principal offers a bonus b to the manager that induces him to choose a success probability of p. The workers observe the bonus b and rationally anticipate that the manager chooses p. What wage does the owner have to pay to workers to make them accept the change? Note that $\Delta x > 0$ implies that $w > w_0$. Thus, the expected utility of workers is given by

$$\begin{aligned} U &= w - x_0 - p\Delta x - \alpha\lambda[w_0 - w]^+ - p\lambda[x_0 + \Delta x - (\alpha x_0 + (1-\alpha)(x_0 + p\Delta x))]^+ \\ &\quad -(1-p)\lambda[x_0 - (\alpha x_0 + (1-\alpha)(x_0 + p\Delta x))]^+ \\ &= w - x_0 - p\Delta x[1 + \lambda(1-(1-\alpha)p)] \\ &= w - x_0 - p(1+\lambda)\Delta x + p^2(1-\alpha)\lambda\Delta x \;\geq\; U_0\ . \end{aligned} \tag{22}$$

The following lemma characterizes the wage payment that workers have to be offered:

Lemma 3. *With probabilistic change the principal has to pay the wage*

$$w = x_0 + p(1+\lambda)\Delta x - p^2(1-\alpha)\lambda\Delta x \; + U_0. \tag{23}$$

boils down to our specification and yields exactly the same results.

[26] As in Sections 2 and 3, the expogenous parameter $1-\alpha$ reflects the weight that the rationally anticipated outcome gets in the formation of the reference point. Expectations are fully rational.

[27] Note that the principal cannot do better by offering a wage payment conditional on the realization of change. Indeed, such a probabilistic wage would even create potential losses with respect to the reference point if success is not realized.

to get workers' consent. This wage function is concave in the probability of change p. It decreases in p if and only if

$$\frac{1+\lambda}{\lambda(1-\alpha)} < 2p. \tag{24}$$

Proof. See the Appendix. □

Equation 24 is more likely to hold if λ is large, α is small, and p is close to 1.[28] If $\alpha = 1$, i.e., if the reference point is fully determined by the status quo, the wage increase is just a linear function of p. The more likely the change, the higher the wage increase that workers demand to accept it. However, if $\alpha < 1$, i.e., if the reference point is partly shaped by rational expectations, the wage increase is a concave and possibly decreasing function of p. The higher the probability of change, the higher the reference point x^r, and the less workers suffer from change. This has important implications for the probability of change that the principal wants to implement. If the weight on expectations for shaping the reference point is sufficiently large, i.e., if α is sufficiently small, the principal will induce change either with probability one or with probability zero. This is shown in the following proposition.

Proposition 7. *The probability of change that the principal wants to implement is characterized as follows:*

(a) If $c < (1-\alpha)\lambda\Delta x$ the principal's problem of inducing the manager to promote change is a convex problem. In this case, the principal will implement a corner solution with $p = 1$ if $\Delta v \geq (1+\alpha\lambda)\Delta x + c$ and $p = 0$ otherwise.

(b) If $c > (1-\alpha)\lambda\Delta x$ the principal's problem is concave. In that case the principal implements $p > 0$ if and only if $\Delta v > (1+\lambda)\Delta x$, in which case p satisfies

$$p = \min\left\{\frac{\Delta v - (1+\lambda)\Delta x}{2[c-(1-\alpha)\lambda\Delta x]}, 1\right\} \tag{25}$$

Proof. See the Appendix □

The proposition shows that if $1-\alpha$ is large, i.e., if rational expectations have a large effect on the reference point, then the manager is induced to choose an extreme solution of

[28] Note that (24) can only hold if $\lambda > 1$.

either $p = 0$ or $p = 1$. Even if an interior solution is optimal, an increase in $1 - \alpha$ increases the probability of change. The reason is that the more the reference point is guided by expectations the less workers suffer from change if it is expected to happen with high probability. The theoretical result that it is optimal to induce the manager to "go all in" resonates with the advice given in the literature on organizational change that if you want to induce change you have to set the expectation that change is coming and that it is unavoidable.

5 Conclusions

This paper shows that reference-dependent preferences can naturally account for several stylized facts about organizational change. First, loss aversion explains why there is often no or slow change in normal times, but a sudden spur in productivity in a crisis. Second, the model shows that when workers have to make concessions in a crisis, these concessions will be about organizational change rather than about cutting wages. This offers a new explanation for the downward rigidity of nominal wages. Third, the model explains why large productivity differences between firms can arise and persist for a long time if firms are founded at different points in time or face idiosyncratic shocks.

Our model has several other interesting implications. For example, it implies that it is more difficult to implement change with older workers. Older workers have a shorter time horizon until they retire. Furthermore, Gächter, Johnson, and Herrmann (2022) report that older people suffer more from loss aversion than younger people, that is, their λ is larger. For both reasons they need to receive a higher compensation to accept change, which makes change more costly to implement.

If the government protects workers with a generous social safety net and if it tries to prevent firm closures, it makes change more difficult to implement. For example, if firms can put their workers in short-term work (subsidized by the government) rather than lay them off, there is less need for workers to make concessions. If there is general unemployment insurance, workers lose less in the case of a crisis and are less willing to accept changes. This resonates with the observation that companies in Anglo-Saxon and Asian countries are often more radical in the implementation of new technologies that require organizational change than their European counterparts.

In the model, we assume that the parameter α, i.e., the weight that the reference point puts on the status quo, is an exogenously given preference parameter. However, preferences may be malleable. A company that wants to implement change could try to shape workers' preferences by reducing α. For example, it could focus attention on the future by making it clear that the world has changed. Alternatively, the company could hire a new manager who has a reputation for pushing change through. All of this shifts attention to the rational expectation that change will happen. If it reduces α, the weight that workers put on the status quo, it reduces the cost to implement change.

The theoretical exploration of these effects and the empirical validation of the impact of reference-dependent preferences on organizational change are interesting and important directions for future research.

References

Akerlof, George A. 1982. "Labor Contracts as Partial Gift Exchange." *The Quarterly Journal of Economics* 97 (4):543.

Alesina, Alberto and Francesco Passarelli. 2019. "Loss Aversion in Politics." *American Journal of Political Science* 63 (4):936–947.

Backus, Matthew. 2020. "Why Is Productivity Correlated With Competition?" *Econometrica* 88 (6):2415–2444.

Balzer, Benjamin, Antonio Rosato, and Jonas von Wangenheim. 2022. "Dutch vs. first-price auctions with expectations-based loss-averse bidders." *Journal of Economic Theory* 205:105545.

Bewley, Truman F. 1999. *Why Wages Don't Fall during a Recession.* Harvard University Press.

Bloom, Nicholas, Erik Brynjolfsson, Lucia Foster, Ron Jarmin, Megha Patnaik, Itay Saporta-Eksten, and John van Reenen. 2017. "What Drives Differences in Management?" *NBER Discussion Paper* 23300, DOI: 10.3386/w23300.

———. 2019. "What Drives Differences in Management Practices?" *American Economic Review* 109 (5):1648–1683.

Bloom, Nicholas, Renata Lemos, Raffaella Sadun, Daniela Scur, and John van Reenen. 2014. "The New Empirical Economics of Management." *Journal of the European Economic Association* 12 (4):835–876.

Branten, Eva, Ana Lamo, and Tairi Room. 2018. "Nominal Wage Rigidity in the EU Countries before and after the Great Recession: Evidence from the WDN surveys." *European Central Bank Working Paper* 2159.

Brown, Alexander L., Taisuke Imai, Ferdinand Vieider, and Colin Camerer. 2023. "Meta-Analysis of Empirical Estimates of Loss Aversion." *Journal of Economic Literature* :forthcoming.

Burke, W. Warner. 2017. *Organization Change: Theory & Practice*. Thousand Oaks, California: SAGE Publications, Inc.

Eliaz, Kfir and Ran Spiegler. 2014. “Reference Dependence and Labor Market Fluctuations.” *NBER Macroeconomics Annual* 28:159–200.

Fallick, Bruce, Daniel Villar, and William Wascher. 2020. “Downward Nominal Wage Rigidity in the United States during and after the Great Recession.” *Federal Reserve Bank of Cleveland Working Paper* 16-02R.

Fehr, Ernst and Gary Charness. 2025. “Social Preferences: Fundamental Characteristics and Economic Consequences.” *Journal of Economic Literature* 63 (forthcoming).

Fehr, Ernst and Lorenz Goette. 2007. “Do Workers Work More if Wages Are High? Evidence from a Randomized Field Experiment.” *American Economic Review* 97 (1):298–317.

Gächter, Simon, Eric J. Johnson, and Andreas Herrmann. 2022. “Individual-Level Loss Aversion in Riskless and Risky Choices.” *Theory and Decision* 92 (3-4):599–624.

Gibbons, Robert and Rebecca Henderson. 2013. “What Do Managers Do?” In *The Handbook of Organizational Economics*, edited by Robert Gibbons and John Roberts. Princeton: Princeton University Press, 680–731.

Grigsby, John, Erik Hurst, and Ahu Yildirmaz. 2021. “Aggregate Nominal Wage Adjustments: New Evidence from Administrative Payroll Data.” *American Economic Review* 111 (2):428–471.

Herweg, Fabian, Heiko Karle, and Daniel Müller. 2018. “Incomplete Contracting, Renegotiation, and Expectation-Based Loss Aversion.” *Journal of Economic Behavior & Organization* 145:176–201.

Holmes, Thomas J. and James A. Schmitz. 2010. “Competition and Productivity: A Review of Evidence.” *Annual Review of Economics* 2 (1):619–642.

Kahneman, Daniel and Amos Tversky. 1979. “Prospect Theory: An Analysis of Decision under Risk.” *Econometrica* 47 (2):263–291.

Karle, Heiko and Heiner Schumacher. 2017. "Advertising and Attachment: Exploiting Loss Aversion through Prepurchase Information." *RAND Journal of Economics* 48 (4):927–948.

Köszegi, Botond and Matthew Rabin. 2006. "A Model of Reference-Dependent Preferences." *Quarterly Journal of Economics* 121 (4):1133–1165.

Kotter, John. 1995. "Leading Change: Why Transformation Efforts Fail." *Harvard Business Review* (March-April):59–67.

Lockwood, Ben, Minh Le, and James Charles Rockey. 2022. "Dynamic Electoral Competition with Voter Loss-Aversion and Imperfect Recall." *SSRN Electronic Journal 4120268* .

Lockwood, Ben and James Rockey. 2020. "Negative Voters? Electoral Competition with Loss-Aversion." *The Economic Journal* 130 (632):2619–2648.

Macera, Rosario. 2018. "Intertemporal Incentives under Loss Aversion." *Journal of Economic Theory* 178:551–594.

Pagel, Michaela. 2016. "Expectations-Based Reference-Dependent Preferences and Asset Pricing." *Journal of the European Economic Association* 14 (2):468–514.

———. 2017. "Expectations-Based Reference-Dependent Life-Cycle Consumption." *Review of Economic Studies* 84 (2):885–934.

Rosato, Antonio. 2023. "Loss Aversion in Sequential Auctions." *Theoretical Economics* 18 (2):561–596.

Schmidt, Klaus M. and Jonas von Wangenheim. 2023. "Organizational Change and Reference-Dependent Preferences." *CRC TRR 190 Discussion Paper* 430, DOI: 10.5282/ubm/epub.107374.

Schmitz, James A. 2007. "What Determines Productivity? Lessons from the Dramatic Recovery of the U.S. and Canadian Iron Ore Industries Following Their Ealry 1980s Crisis." *Journal of Political Economy* 113 (3):582–625.

Syverson, Chad. 2004. "Product Substitutability and Productivity Dispersion." *Review of Economics and Statistics* 86 (2):534–550.

———. 2011. “What Determines Productivity?” *Journal of Economic Literature* 49 (2):326–365.

Van Bilsen, Servaas, Roger JA Laeven, and Theo E Nijman. 2020. “Consumption and Portfolio Choice under Loss Aversion and Endogenous Updating of the Reference Level.” *Management Science* 66 (9):3927–3955.

von Wangenheim, Jonas. 2021. “English versus Vickrey Auctions with Loss-Averse Bidders.” *Journal of Economic Theory* 197:105328.