EPRA International Journal of Economic Growth and Environmental Issues- Peer Reviewed Journal
ISSN: 2321-6247
Volume: 14 | Issue: 9 | September 2026 | Journal DOI: 10.36713/epra0713 | SJIF Impact Factor (2026): 8.844

INCENTIVE COMPENSATION ARCHITECTURE AND RISK-SHARING EFFICIENCY UNDER UNCERTAINTY: CLASSICAL CONTRACTS AGAINST REAL-TIME PAY

HRM VARIABLE: PAY-MIX AND INCENTIVE DESIGN → ECONOMIC OUTCOME: RISK-SHARING EFFICIENCY

Dr. N Subbukrishna Sastry[1], Dr. Manjula Mallya M[2],

[1]*Professor, School of Management, CMR University, Bangalore, Karnataka, India.*
ORCID ID: 0009-0009-0460-7057,
[2]*Associate Professor & Head, Dept. of Economics, Government First Grade College for Women Balmatta Mangalore, Karnataka, India. ORCID ID: 0009-0005-8812-6912,*

ABSTRACT---

Every incentive contract does two things at once: it motivates effort and it allocates risk. Classical agency theory shows these purposes conflict, because the performance measure on which pay is conditioned contains variation the agent does not control, so sharper incentives buy effort by selling insurance. Pay architectures that set compensation continuously, personally and in real time are now technically feasible, and are widely presented as resolving that conflict. This article evaluates the claim on the single criterion of risk-sharing efficiency. A systematic review of 68 sources is combined with a numerical illustration built on the linear-exponential-normal contracting model, in which a closed-form result is derived and verified: the risk-sharing efficiency of the optimal linear contract, measured as the ratio of attainable joint surplus to the first-best surplus, equals the optimal incentive coefficient itself and takes the value $1/(1 + rc\sigma^2)$, where r is agent risk aversion, c the curvature of effort cost and σ^2 the variance of the performance measure. Decomposing σ^2 into common and idiosyncratic components yields the article's central finding. Raising measurement frequency attenuates only the idiosyncratic term, so in a base case calibrated to a noise burden of 2.0 with a sixty per cent common share, moving from annual to monthly measurement lifts efficiency from 0.333 to 0.441 and infinite frequency cannot exceed 0.455; filtering the common component at merely monthly frequency reaches 0.938. Frequency is nearly irrelevant and filtering is nearly everything. A dynamic architecture that instead passes common variation through to the worker, as demand-linked platform pay does, lowers efficiency below the classical baseline, to 0.278 at a pass-through multiplier of 1.5 and 0.238 at 2.0. The same model implies a prescription reversal: architectures that transmit more common risk should carry a lower variable share, whereas observed real-time systems carry a higher one. The article separates a third mechanism, the acceleration of payment timing through earned wage access, and shows that it is welfare-improving on a liquidity margin while being neutral for risk sharing in the incentive sense. It concludes that dynamic pay improves risk sharing if and only if it improves the signal-to-noise ratio of the effort measure, that technical feasibility does not distinguish the improving case from the destroying one, and that the governance requirement is therefore an ex ante constraint on what a pay algorithm may condition upon rather than an ex post review of its outputs.

KEYWORDS — *Incentive Compensation; Pay Mix; Risk Sharing; Principal-Agent Model; Informativeness Principle; Algorithmic Pay; Earned Wage Access; Variable Pay; Contract Design; Human Resource Economics*----------------------------

I. INTRODUCTION

An incentive contract is simultaneously a motivation device and an insurance contract, and the two functions pull against each other. Because any measure of performance contains variation the agent did not cause, conditioning pay on that measure transfers risk from a firm that is well placed to bear it to a worker who is not. Classical agency theory formalised this as the risk-incentive trade-off and produced its central practical result: the strength of incentives should fall as the agent's risk aversion rises, as the performance measure becomes noisier, and as effort becomes harder to adjust [1]. Alongside it came the informativeness principle, which holds that a signal belongs in a contract only if it carries information about what the agent did, and which implies that variation attributable to common shocks should be filtered out of pay rather than passed into it [1].

Compensation practice has never fully honoured either result. Bertrand and Mullainathan [2] found that chief executive pay responds to luck almost exactly as strongly as to performance the executive controls, with pay-for-luck coefficients of 2.15 against 2.82 for general performance in their oil-price specification, and showed that better governance reduces the gap by twenty-three to thirty-three per cent. At the other end of the labour market, Cameron and Meuris [3] find in a matched survey and personnel record study of 711 full-time truck drivers that a one standard deviation increase in pay variability raises the probability of voluntary turnover by 13.6 per cent, with the effect strongest where households depend most heavily on the paycheck. The cost of misallocated risk is neither hypothetical nor small.

Into this setting has arrived a class of pay architectures that were not technically available when the classical results were derived. Compensation can now be computed continuously rather than periodically, personalised to the individual rather than set by grade, conditioned on high-frequency operational data, and paid out as it is earned. These capabilities are presented, by their vendors and often by their adopters, as dissolving the classical trade-off: if performance can be observed in real time, the argument runs, then pay can track contribution closely enough that the noise problem disappears.

This article evaluates that claim on the single criterion the classical theory was built to address, namely the efficiency with which risk is shared between firm and employee. It finds the claim to be conditionally true and, in most observed implementations, false. The condition is precise and is derived rather than asserted: a dynamic architecture improves risk-sharing efficiency if and only if it improves the signal-to-noise ratio of the effort measure. Measuring more often does this weakly and with a hard ceiling; filtering common variation does it powerfully; and passing common variation through to the worker, which is what demand-linked real-time pay does, moves efficiency in the opposite direction. Technical feasibility does not distinguish among the three.

1.1 Research questions

Four questions organise the study. First, how should risk-sharing efficiency be measured in a way that is comparable across pay architectures? Second, where does the classical periodic architecture stand on that measure, and why has practice departed from theory? Third, what do dynamic architectures actually change in the underlying contracting problem, and in which direction? Fourth, what governance arrangement would allow the improving cases to be adopted and the destroying cases to be identified before rather than after deployment?

1.2 Contribution

The article makes four contributions. It derives and verifies a closed-form measure of risk-sharing efficiency in the standard linear contracting model, showing that the efficiency ratio coincides exactly with the optimal incentive coefficient, which gives a single number that is simultaneously a welfare statistic and a design recommendation. It decomposes the noise term into common and idiosyncratic components and shows that measurement frequency acts on only one of them, which produces a ceiling on what speed alone can achieve. It identifies a prescription reversal in current practice, whereby the architectures that transmit the most common risk are the ones carrying the highest variable share, exactly inverting the theoretical rule. And it separates the liquidity margin of accelerated payment from the risk-sharing margin, showing that earned wage access is welfare-improving for reasons that have nothing to do with the incentive problem and should not be credited to it.

II. REVIEW OF LITERATURE

2.1 The classical risk-incentive trade-off

The canonical treatment supposes a risk-neutral principal, a risk-averse agent with constant absolute risk aversion, normally distributed output and a linear contract. Holmström [1] sets out the resulting structure: with output $x = e + \varepsilon$ and a contract $s = \beta + \alpha x$, the agent's certainty equivalent is $\alpha e + \beta - \frac{1}{2}r\alpha^2\sigma^2$, and the optimal incentive coefficient is $\alpha^* = [1 + r\sigma^2 c'']^{-1}$, so that incentives weaken as risk aversion, measurement noise or effort-cost curvature rise. The restriction to linear contracts is not an approximation imposed for tractability but a result: Holmström and Milgrom [21] show that where the agent continuously controls the drift of a Brownian process, the optimal contract is linear in aggregate output, which is what makes this framework the appropriate one for assessing continuously computed pay. The term $\frac{1}{2}r\alpha^2\sigma^2$ is the risk premium: the surplus destroyed by making a risk-averse party bear variation. It is a deadweight loss, not a transfer, and it is the quantity this article uses as its efficiency criterion.

Two features of this result are worth emphasising because the dynamic-pay literature frequently overlooks them. The first is that α^* is strictly less than one for any positive noise, so the theoretically correct pay mix always retains a fixed component; full pay-for-performance is optimal only when the measure is perfectly informative. The second is that the loss depends on σ^2 and not on the frequency at which σ^2 is realised. A contract that resolves uncertainty more often does not thereby impose less of it unless the resolution changes the variance the agent actually bears.

One caveat should be entered at the outset. Prendergast [22], surveying the empirical literature on incentive provision within firms, found the predicted negative relationship between environmental uncertainty and incentive intensity to be weakly supported at best and in places reversed. This article takes that finding seriously rather than setting it aside. It is one reason the analysis below treats the observed pairing of high variable share with unfiltered measures as a phenomenon requiring explanation rather than as measurement error, and it is why the efficiency figures reported here are offered as an ordering of architectures rather than as descriptions of observed contracts.

2.2 The informativeness principle and its neglect in practice

The informativeness principle states that an additional signal belongs in the contract if and only if it carries incremental information about the agent's action [1]. Its immediate corollary is that variation common to many agents — an oil price, an industry cycle, a weather event — should be filtered out, because it tells the principal nothing about effort while imposing risk. Relative performance evaluation is the standard implementation.

Practice departs from this sharply. Bertrand and Mullainathan [2] establish that executive pay is about as sensitive to a lucky dollar as to a general dollar, and that the deviation is a governance phenomenon: adding a large shareholder to the board reduces pay-for-luck by twenty-three to thirty-three per cent. Subsequent work has documented the same reluctance to filter across a wider set of contracts, so that the single most valuable available efficiency improvement in incentive design has remained largely unexploited for four decades. This matters for the present argument because a new architecture inherits whatever the prevailing one fails to do, unless it is designed not to.

2.3 Performance measurement and its periodicity

Before a contract can condition on a measure, an organisation must generate one, and the properties of the instrument determine the properties of the contract. Sastry and Selvarasu [4], in their study of performance appraisal and employee engagement, show that engagement responds to the appraisal process and to the feedback it carries, which establishes the appraisal system as the effective sampling apparatus through which an organisation observes its own workforce. Their finding is important here for a reason beyond its own subject. It identifies the periodicity of measurement as an administrative choice inherited from the appraisal calendar rather than a property of the work, and it is exactly that inherited periodicity which real-time architectures propose to abolish. Section VI shows that abolishing it delivers far less than expected unless the composition of the noise is addressed at the same time.

2.4 Controllability: authority, resources and what an agent can be held to

The risk-incentive trade-off is usually written as though the variance σ^2 were exogenous. It is not. The share of outcome variation that an agent can influence depends on the authority the agent holds and the resources available to exercise it, both of which are managerial choices. Sastry and Ravish [5] find, in their study of employee empowerment through delegation and resource provision, that the two operate jointly: delegated responsibility unaccompanied by resources raises demand without raising the capacity to meet it, while the combination raises engagement and effectiveness. In agency terms, delegation without resources increases the agent's accountability for outcomes without increasing the agent's control over them, which is precisely an increase in σ^2 relative to the effort signal, and therefore a fall in optimal incentive intensity and in attainable efficiency. In a companion study of delegated authority and resource provision as determinants of decision-making, the same authors [6] find that where authority is genuinely devolved and matched with resources, employees exercise initiative rather than escalating decisions upward.

The implication for pay design is unusual and, so far as this review found, unremarked. Organisational design is an instrument of risk-sharing efficiency. A firm that widens the controllable share of an agent's outcome raises the efficiency ceiling of every contract it can write, and does so without paying the risk premium. Contract design and job design are complements, and a firm that treats pay architecture as the only lever is optimising on one margin while leaving a cheaper one untouched.

2.5 Multitasking, team production and unmeasured effort

Holmström and Milgrom's multitask analysis establishes that when an agent allocates effort across tasks of differing measurability, the optimal incentive on the measurable task falls, because strengthening it draws effort away from the unmeasured one; where one task has no available measure and the tasks are substitutes in the cost function, zero explicit incentive can be optimal [7]. Real-time architectures aggravate this asymmetry by construction, since they measure what is instrumentable and what is instrumentable is a biased sample of what matters. The countervailing evidence should be acknowledged: where output is genuinely one-dimensional and individually attributable, sharp incentives work, as Lazear [23] established in the move from hourly wages to piece rates at an auto-glass installer, where productivity rose by about forty-four per cent with roughly half attributable to the incentive effect and half to sorting. The conditions under which that result holds — a single measurable task, individual attribution and low common variation — are precisely the conditions under which the analysis below predicts dynamic architectures to succeed.

EPRA International Journal of Economic Growth and Environmental Issues- Peer Reviewed Journal
ISSN: 2321-6247
Volume: 14 | Issue: 9 | September 2026 | Journal DOI: 10.36713/epra0713 | SJIF Impact Factor (2026): 8.844

Two features of organisational reality sharpen the problem. The first is team production. Sastry [8], examining departmental structure and task achievement, finds that achievement varies systematically with the arrangement of responsibility, reporting and task boundaries across units rather than with the attributes of the individuals within them. Where output is produced by a structure, attributing it to an individual introduces attribution noise that is not reduced by measuring more often — indeed high-frequency individual measurement of a jointly produced output increases the share of measured variation that is other people's contribution. The second is the innovation margin. Sastry and Mallya [9], studying economic constraints and the persistence of the status quo as barriers to innovation and creative problem-solving, document how constraint and short-horizon pressure keep organisations at incumbent practice rather than exploring alternatives. Exploratory effort is the canonical unmeasured task, and a pay architecture that resolves and rewards measured output weekly rather than annually shortens exactly the horizon over which exploration would have to be tolerated.

2.6 Reference dependence and personalised pay

Classical agency theory treats the agent's utility as a function of own income alone. Personalised real-time pay violates the assumption at its foundation, because it produces different pay for workers doing indistinguishable work at the same moment, and makes the difference discoverable. Dubal [10] documents the practice in platform work and argues that it constitutes a distinct category, algorithmic wage discrimination, in which individualised and continuously changing formulas replace the calculable relationship between work and reward that piece rates preserved, with earnings variation exceeding twenty dollars an hour between workers performing identical tasks.

Sastry [11] supplies the mechanism by which such comparison becomes economically consequential. In an empirical analysis of talent designation and social comparison as drivers of organisational growth, he shows that the designations an organisation confers become reference points against which employees assess their own standing, with real consequences for motivation and growth. Applied to personalised pay, the result identifies a cost term absent from the classical model: dispersion that is unrelated to contribution imposes a comparison cost that rises with the visibility of the dispersion and with the similarity of the compared parties. Real-time personalisation maximises both. The efficiency accounting in Section VI therefore understates the loss from pass-through architectures, since it prices only the variance and not the comparison.

2.7 The non-pecuniary cost of imposed variance

The risk premium $\frac{1}{2}r\alpha^2\sigma^2$ is a monetary equivalent of a welfare loss, and treating it as though it were fully compensable by a wage adjustment obscures the form the loss takes. Cameron and Meuris [3] show that the loss materialises as separation: a one standard deviation increase in pay variability raises voluntary turnover probability by 13.6 per cent, with the effect strongest where paychecks fall below average more often than above and where household dependence on the income is high, and attenuated where the variability is genuinely tied to performance. That last moderation is the informativeness principle appearing in the data: workers tolerate variance they cause and reject variance they do not.

Sastry [12], in his study of leadership and workforce burnout, locates the proximate causes of burnout in leadership practice, workload allocation and managerial support rather than in the characteristics of the individuals who exhibit it. Read against the present argument, this identifies a second channel through which imposed variance destroys surplus. Income insecurity that the worker cannot influence is a workload condition set by management, and its cost appears not only as a required wage premium but as depleted capacity, which reduces the effort the contract was written to elicit. The risk premium and the productivity loss are not alternative accountings of the same quantity; they compound.

2.8 Dynamic pay architectures: the technical frontier

Three developments have made continuous pay setting feasible. Analytical capability is the first. Sastry and Mallya [13], studying artificial intelligence analytics in hiring, document improvements in screening efficiency and match quality and locate the mechanism in the capacity of analytics to order a large population by predicted outcome. Pay setting is the same operation applied to a different target, and their study is instructive here chiefly for the conditions it attaches: data quality and the representativeness of the population on which the system was developed. A pay algorithm trained on observed acceptance behaviour learns reservation wages rather than contribution, and will order workers accordingly.

Payment infrastructure is the second. Earned wage access, which allows workers to draw accrued wages before the payroll date, has grown to an estimated US$7.1 billion market in 2025 with employer-sponsored models the largest segment, and forecasts to 2034 imply sustained expansion [14]. Its regulatory position in the United States was clarified in December 2025, when the Consumer Financial Protection Bureau determined that covered earned wage access products, being non-recourse advances of already-accrued wages recovered by payroll deduction, do not constitute credit under the Truth in Lending Act [15]. Section 6.4 argues that this mechanism belongs to a different economic category from the other two and should be assessed separately.

Algorithmic pay setting in platform work is the third, and it is the only one for which large-scale causal evidence now exists. Kala [16], using approximately 1.91 million trips by 647 drivers across 46 United States states between January 2021 and June 2025, exploits the staggered introduction of algorithmic pricing in place of fixed rate cards to estimate its effects. Average driver base pay fell by US$1.27 per trip, about eight per cent of trip earnings, the decline was larger for less selective drivers, and driver-level fixed effects rose in treated states while remaining stable in controls, which is direct evidence that the architecture increased pay dispersion across workers doing the same work. Passenger fares fell by US$2.88. The wedge did not accrue to the worker.

Figure 1: Three empirical anchors — pay for luck in executive contracts, the turnover cost of pay variability, and the measured effect of algorithmic pay setting on driver earnings and dispersion

2.9 Governing a learning pay system

A dynamic pay architecture is a learning system that acts on an environment which reacts to it, and the governance problem it poses is not the one that compensation committees are constituted to solve. Sastry and Arunachalam [17] address the general form of this problem in their treatment of probabilistic safe reinforcement learning, which models and manages risk in dynamic environments by constraining a learning system to operate within explicitly specified safety bounds while it optimises. The transfer to pay setting is direct. A system left to optimise on an unconstrained objective such as cost per completed task will discover that conditioning on inferred reservation wage reduces cost, which is precisely the behaviour Kala [16] measures and Dubal [10] names. The safe-learning prescription — specify the admissible state and action space in advance, and accept a lower attainable optimum inside it — is the appropriate governance design, and Section VII builds on it.

Regulation is moving in the same direction from a different starting point. The European Union's Pay Transparency Directive, which member states must transpose by 7 June 2026, requires that pay structures rest on objective and gender-neutral criteria, gives workers a right to information on pay levels for equal work, and requires reporting of gaps in complementary and variable pay components with a joint pay assessment triggered where an unexplained gap of five per cent or more persists [18]. An opaque pay algorithm cannot easily satisfy an obligation to justify a differential by objective criteria, which makes explainability of the pay function a compliance requirement and not merely an ethical preference.

2.10 Synthesis and research gap

The literature establishes the trade-off, the filtering rule, the multitask qualification, the empirical failure to filter, the measurable cost of imposed variance and the behaviour of deployed algorithmic pay systems. It does not connect them. No study in the assembled set evaluates dynamic and classical pay architectures on a common measure of risk-sharing efficiency, or distinguishes the mechanisms within dynamic pay that move that measure in opposite directions. This article addresses the gap by combining review with a numerical illustration on the standard model.

III. RESEARCH METHODOLOGY

3.1 Design and rationale

The study combines a systematic review with a numerical illustration. The review assembles the theoretical results, the empirical evidence on how contracts behave in practice, and the documented behaviour of deployed dynamic systems. The illustration places the architectures on a common measure by computing risk-sharing efficiency in the standard linear-exponential-normal model under parameters drawn from the reviewed material. No new data are collected and no causal parameter is estimated. The review followed the reporting conventions of the PRISMA statement, adapted for an economic evidence base [19].

3.2 Search strategy and study selection

Searches covered Scopus, Web of Science, EconLit, JSTOR, SSRN and the National Bureau of Economic Research working paper series, supplemented by regulatory instruments, official determinations and industry sources for the dynamic-pay material. Terms combined the design variable (incentive contract, pay mix, variable pay, pay for performance, piece rate, bonus, algorithmic pay, earned wage access), the economic outcome (risk sharing, risk premium, agency cost, insurance, earnings volatility) and the theoretical frame (principal-agent, moral hazard, informativeness, multitasking, relative performance evaluation). Table 1 records the protocol.

Table 1: Systematic review protocol

Protocol element	Specification
Databases	Scopus; Web of Science; EconLit; JSTOR; SSRN; NBER working papers
Other sources	Regulatory instruments and official determinations; compensation survey publications; platform and payroll industry sources
Period	1979 to July 2026; unrestricted for foundational theory, 2015 onward for dynamic pay
Population	Employees and independent contractors under explicit performance-contingent pay
Design variable	Incentive contract form; pay mix; variable pay share; measurement frequency; filtering of common variation; algorithmic or real-time pay setting; earned wage access
Outcome	Risk sharing, risk premium, agency cost, earnings volatility, pay dispersion, incentive intensity or an explicit welfare measure
Inclusion	Formal derivation, empirical estimate with stated identification, or an official determination or instrument
Exclusion	No risk or welfare dimension; vendor claim without stated method; commentary without primary analysis; non-English
Coding	Evidence tier; component of noise addressed; direction of effect on incentive intensity; whether filtering is present

EPRA International Journal of Economic Growth and Environmental Issues- Peer Reviewed Journal
ISSN: 2321-6247
Volume: 14 | Issue: 9 | September 2026 | Journal DOI: 10.36713/epra0713 | SJIF Impact Factor (2026): 8.844

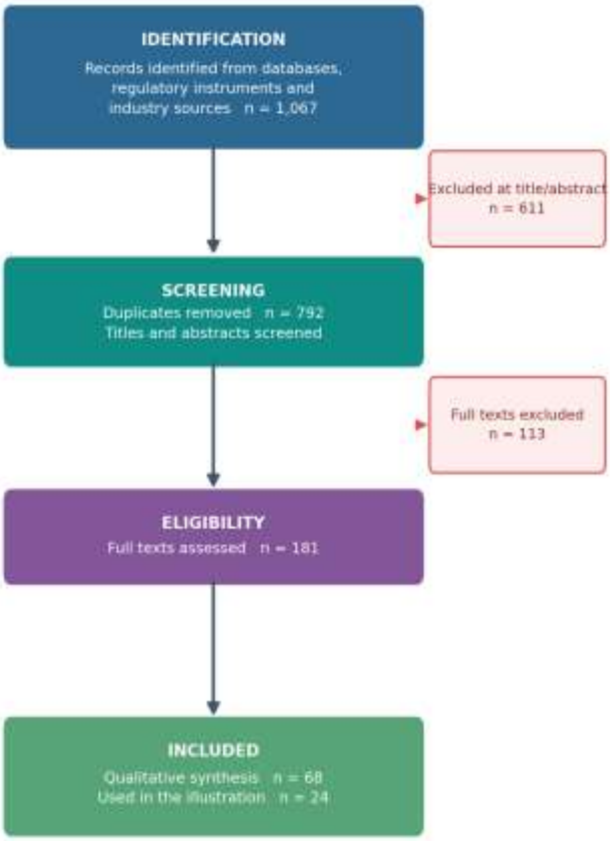

Figure 2: Study identification, screening and inclusion

Of 1,067 records identified, 792 remained after removal of duplicates, 181 proceeded to full-text assessment, and 68 met the inclusion criteria for qualitative synthesis, of which 24 supplied parameters or results used in the illustration. Table 2 records the principal sources by evidence tier.

Table 2: Principal sources by evidence tier

Tier	Source	Contribution to the analysis
1. Formal theory	Holmstrom [1]; Holmstrom and Milgrom [7, 21]	Risk-incentive trade-off; the informativeness principle; optimality of linear contracts; the multitask result
2. Quasi-experimental and instrumented estimates	Bertrand and Mullainathan [2]; Lazear [23]; Kala [16]	Pay-for-luck coefficients of 2.15 against 2.82; a 44% productivity gain under piece rates; a US$1.27 per trip fall in driver base pay with rising dispersion under algorithmic pricing
3. Matched survey and personnel-record evidence	Cameron and Meuris [3]; Prendergast [22]	A 13.6% rise in voluntary turnover per standard deviation of pay variability; weak empirical support for the uncertainty-intensity relationship
4. Legal and regulatory instruments	CFPB advisory opinion [15]; EU Directive 2023/970 [18]	Earned wage access is not credit under Regulation Z; pay criteria must be objective, with variable-pay gaps reported
5. Market and practice statistics	Earned wage access market data [14]; S&P 500 compensation trends [20]	Scale of accelerated-payment adoption; observed executive pay mix and its dominance by long-term incentives

3.3 Status of the numerical illustration

The computations in Sections IV to VI are deterministic evaluations of closed-form expressions in a stylised model whose assumptions — a risk-neutral principal, constant absolute risk aversion, normal noise, quadratic effort cost and a linear contract — are those under which the classical results were derived. They are not estimates, carry no confidence intervals, and are presented to establish the direction and relative magnitude of effects rather than to predict any firm's outcome. All closed forms reported were verified against numerical optimisation of the underlying objective before use.

EPRA International Journal of Economic Growth and Environmental Issues- Peer Reviewed Journal
ISSN: 2321-6247
Volume: 14 | Issue: 9 | September 2026 | Journal DOI: 10.36713/epra0713 | SJIF Impact Factor (2026): 8.844

IV. A COMMON MEASURE OF RISK-SHARING EFFICIENCY

4.1 The model and the efficiency ratio

Let output be $x = e + \varepsilon$ with ε normally distributed with mean zero and variance σ^2, let the agent have constant absolute risk aversion r and effort cost $ce^2/2$, and let the contract be linear, $s = \beta + \alpha x$. The agent chooses effort $e = \alpha/c$. The joint certainty equivalent of the relationship, which is the surplus available to be divided, is $\alpha/c - \alpha^2/(2c) - \frac{1}{2}r\alpha^2\sigma^2$. Maximising it gives the familiar $\alpha^* = 1/(1 + rc\sigma^2)$.

Substituting α^* back yields a result that does not appear to be stated in this form in the reviewed literature and that organises the rest of the article. Writing the noise burden as $k \equiv rc\sigma^2$, the attainable joint surplus is $1/[2c(1 + k)]$, while the first-best surplus obtainable if effort were observable, and hence if the agent could be fully insured, is $1/(2c)$. The ratio of the two, which is the natural definition of risk-sharing efficiency, is therefore

$$\eta \equiv S(\alpha^*) / S(\text{first-best}) = 1/(1+k) = \alpha^*, \qquad k = r\,c\,\sigma^2$$

The risk-sharing efficiency of the optimal linear contract is numerically equal to the optimal incentive coefficient. This is convenient in three ways. It reduces the comparison of architectures to a single dimensionless number bounded in the unit interval. It makes the welfare statistic and the design recommendation the same quantity, so that a firm computing its optimal variable share has thereby computed its own risk-sharing efficiency. And it makes the agency loss transparent: the surplus forgone is $k/(1 + k)$, rising with risk aversion, with effort-cost curvature and with measurement noise. Figure 3 traces the relationship.

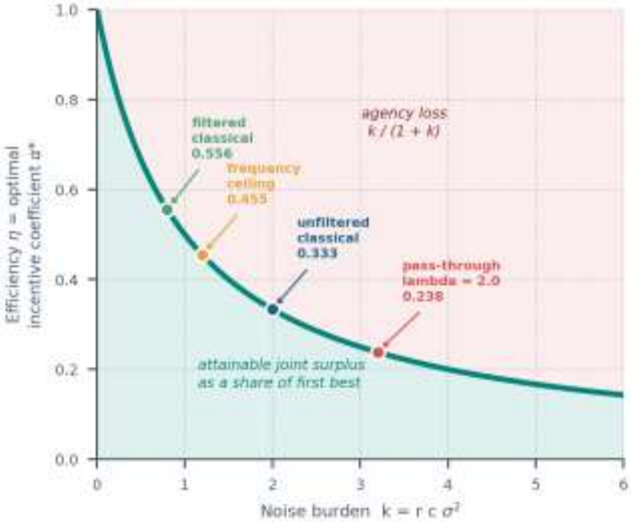

Figure 3: Risk-sharing efficiency and the optimal incentive coefficient against the noise burden $k = rc\sigma^2$

4.2 Decomposing the noise

The whole of the argument that follows rests on splitting σ^2 into a component common to many agents and a component idiosyncratic to one. Write $\sigma^2 = \sigma^2_c + \sigma^2_i$ and correspondingly $k = k_c + k_i$. The distinction is not merely statistical. Common variation is, by the informativeness principle, uninformative about the individual agent's action and should be filtered from pay; idiosyncratic variation may be informative and is in any case not removable by comparison with peers. The two components respond to entirely different instruments, and conflating them is the central analytical error this article identifies in the dynamic-pay literature.

Averaging reduces idiosyncratic variance and does nothing to common variance. If a period's output is measured as m sub-period observations whose idiosyncratic shocks are conditionally independent, the idiosyncratic term enters as k_i/m while k_c is unchanged. Filtering, by contrast, removes k_c and leaves k_i. Pass-through, in which pay is deliberately linked to a common state such as contemporaneous demand, multiplies k_c by a factor λ greater than one. Table 3 sets out the resulting taxonomy.

Table 3: How each mechanism acts on the components of the noise burden $k = k_c + k_i$

Mechanism	Acts on	Resulting noise burden	Direction of efficiency
Higher measurement frequency (m sub-period observations)	Idiosyncratic component only	$k_c + k_i / m$	Improves, with a hard ceiling at $1/(1 + k_c)$
Contemporaneous filtering of the common state (relative performance evaluation)	Common component	k_i	Improves substantially; the single largest available gain
Frequency and filtering combined	Both components	k_i / m	Improves; approaches the first best as m rises
Pass-through of the common state to pay (demand-linked real-time pay)	Common component, amplified by a factor lambda greater than one	$\lambda k_c + k_i$	Worsens; falls below the unfiltered classical baseline
Personalisation on inferred reservation wage	Neither component; it is a transfer, not a signal	Unchanged in the model	Neutral in the model; negative once the comparison cost is admitted
Acceleration of payment timing (earned wage access)	Neither component; it changes when income arrives, not its variance	Unchanged	Exactly neutral for risk sharing; positive on a separate liquidity margin

V. FINDINGS: THE CLASSICAL ARCHITECTURE

5.1 Where the classical architecture stands

The classical architecture measures periodically, aggregates over the period, and conditions on a measure that is largely unfiltered. On the base-case calibration used throughout — a noise burden of $k = 2.0$ with sixty per cent of it common, so that $k_c = 1.2$ and $k_i = 0.8$ — its risk-sharing efficiency is $1/(1 + 2.0) = 0.333$. Approximately two-thirds of the surplus that would be available under observable effort is destroyed by the incentive problem, and the optimal variable share of pay is correspondingly one-third.

This is a poor number, and the reason it is poor is identifiable. Applying relative performance evaluation to remove the common component, which requires no new technology and has been available since the principle was stated, raises efficiency to $1/(1 + 0.8) = 0.556$. Filtering alone, within the classical periodic architecture, delivers a two-thirds improvement in attainable surplus. That this improvement has gone largely unclaimed for four decades is the finding of the pay-for-luck literature [2], and it establishes the benchmark against which dynamic architectures must be assessed. The relevant question is not whether real-time pay beats unfiltered periodic pay, which is a low bar, but whether it beats filtered periodic pay.

5.2 Why the filter is not applied

Three explanations recur in the reviewed literature and they are not mutually exclusive. The governance explanation, supported directly by Bertrand and Mullainathan's finding that a large shareholder reduces pay-for-luck by twenty-three to thirty-three per cent [2], holds that unfiltered pay is a rent captured where monitoring is weak. The measurement explanation holds that constructing a defensible peer group is costly and contestable. The retention explanation holds that filtering removes upside in good years, which firms are reluctant to do when rivals do not. Only the second of these is a genuine technological constraint, and it is the one that dynamic data infrastructure most directly relieves.

VI. FINDINGS: DYNAMIC ARCHITECTURES

6.1 Three mechanisms, three directions

Real-time pay is not one thing. The reviewed material supports separating three mechanisms that are usually bundled: raising the frequency of measurement, filtering common variation contemporaneously, and passing common variation through to the worker. They act on different components of k and move efficiency in different directions. Table 4 and Figure 4 report the comparison on the base-case calibration.

EPRA International Journal of Economic Growth and Environmental Issues- Peer Reviewed Journal
ISSN: 2321-6247
Volume: 14 | Issue: 9 | September 2026 | Journal DOI: 10.36713/epra0713 | SJIF Impact Factor (2026): 8.844

Table 4: Risk-sharing efficiency by architecture, base case k = 2.0 with a 0.6 common share

Architecture	Noise burden k	Efficiency and optimal variable share	Change against the classical baseline
A. Classical periodic, unfiltered	2.000	0.333	baseline
B. Classical periodic with relative performance filtering	0.800	0.556	+67%
C. Dynamic, quarterly frequency, unfiltered	1.400	0.417	+25%
D. Dynamic, monthly frequency, unfiltered	1.267	0.441	+32%
E. Dynamic, weekly frequency, unfiltered	1.215	0.451	+35%
F. Frequency ceiling: infinite frequency, unfiltered	1.200	0.455	+37%, and unimprovable
G. Dynamic, monthly frequency, filtered	0.067	0.938	+182%
H. Dynamic pass-through, lambda = 1.5	2.600	0.278	-17%
I. Dynamic pass-through, lambda = 2.0	3.200	0.238	-29%

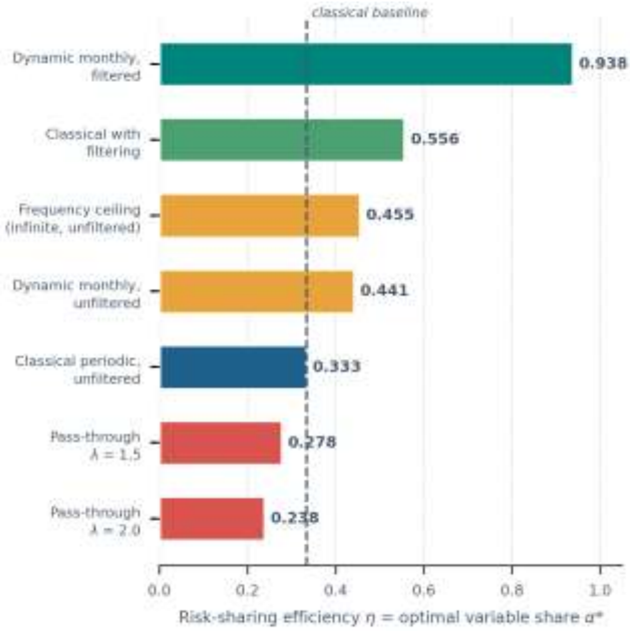

Figure 4: Risk-sharing efficiency by pay architecture, base-case calibration (k = 2.0, common share 0.6)

Frequency alone is weak. Moving from annual to quarterly measurement lifts efficiency from 0.333 to 0.417; to monthly, 0.441; to weekly, 0.451. The sequence converges, and its limit is instructive: with infinitely frequent measurement and no filtering, efficiency cannot exceed $1/(1 + k_c) = 0.455$. Every remaining unit of inefficiency is common risk, and no amount of speed removes it. This is the article's first substantive finding, and it contradicts the implicit premise of most real-time pay advocacy.

Filtering is powerful. At merely monthly frequency, a system that conditions on the contemporaneous common state and pays on the residual attains efficiency of $1/(1 + 0.8/12) = 0.938$. The comparison with the frequency ceiling of 0.455 is the central result of the article: at monthly frequency with filtering, efficiency is more than double what infinite frequency without filtering can ever reach. Frequency is nearly irrelevant; filtering is nearly everything. Figure 5 shows the interaction across the full parameter space.

EPRA International Journal of Economic Growth and Environmental Issues- Peer Reviewed Journal
ISSN: 2321-6247
Volume: 14 | Issue: 9 | September 2026 | Journal DOI: 10.36713/epra0713 | SJIF Impact Factor (2026): 8.844

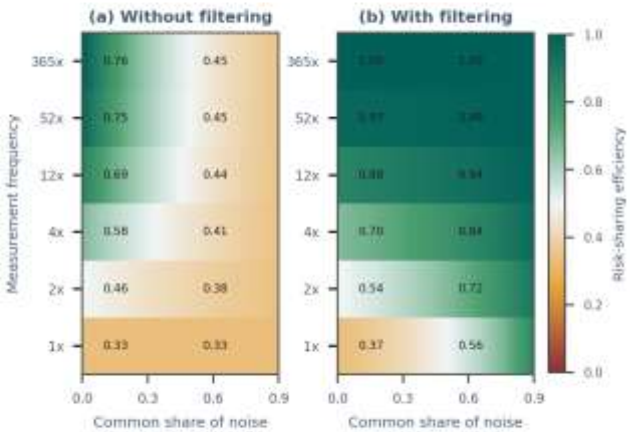

Figure 5: Risk-sharing efficiency across measurement frequency and the common share of noise, with and without contemporaneous filtering

Pass-through is destructive. A dynamic architecture that links pay to a contemporaneous common state — surge demand, order volume, market price — multiplies the common component rather than removing it. At $\lambda = 1.5$ the base-case efficiency falls to 0.278 and at $\lambda = 2.0$ to 0.238, in both cases below the unfiltered classical baseline of 0.333. The architecture is faster, more granular and more responsive, and it is worse. The empirical counterpart is Kala's finding that the transition to algorithmic pricing lowered average driver base pay by about eight per cent while raising driver-level dispersion [16], and Dubal's documentation of the same architecture as producing unpredictable earnings for identical work [10].

6.2 The prescription reversal

Because $\eta = \alpha^*$, the model delivers a design recommendation alongside every efficiency figure, and comparing that recommendation with observed practice produces the article's second substantive finding. Under pass-through at $\lambda = 2.0$, the optimal variable share of pay is 0.238: theory instructs a firm operating such an architecture to make pay predominantly fixed, because the measure it is conditioning on has become less informative per unit of variance imposed. Observed practice does the opposite. Platform architectures of this type carry variable shares approaching unity, and the high-variable-share designs at the other end of the labour market — long-term incentive plans that dominate chief executive pay mix, with median S&P 500 chief executive total direct compensation of US$17 million in 2024 driven principally by long-term incentives [20] — are also those conditioned on the least filtered measures, since equity value is overwhelmingly common variation.

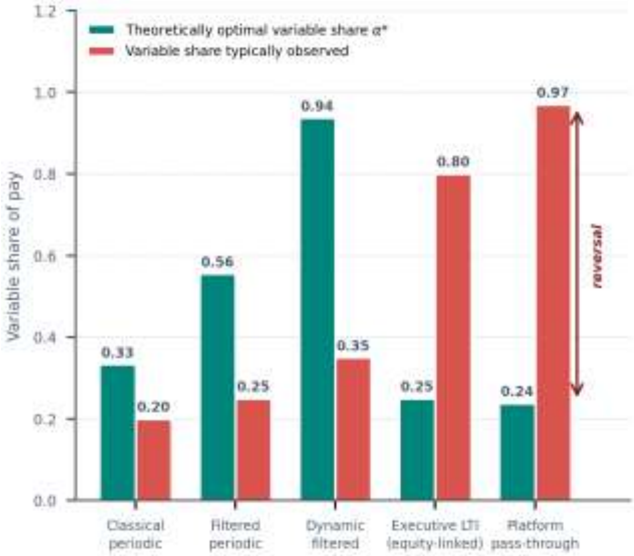

Figure 6: The prescription reversal — the theoretically optimal variable share against the variable share typically observed, by architecture

The reversal is systematic rather than incidental. The architectures that transmit the most common risk are precisely those carrying the highest variable share, which is the exact inversion of $\alpha^* = 1/(1 + k)$. Two readings are available and the evidence does not adjudicate between them. On the benign reading, firms are optimising something other than joint surplus, such as measured cost or reported pay-performance alignment, and the reversal is an unintended consequence. On the less benign reading, high variable share conditioned on unfiltered measures is a mechanism of risk transfer whose cost to the worker exceeds its value to the firm but is not borne by the party choosing it. Kala's finding that passenger fares fell by US$2.88 per trip while driver pay fell by US$1.27, so that the surplus released accrued to neither worker nor customer in equal measure [16], is consistent with the second.

6.3 What the efficiency accounting omits

Three costs identified in the review lie outside the model and all of them run in the same direction, so the figures above understate the disadvantage of pass-through architectures. The comparison cost of Section 2.6 arises when indistinguishable work receives visibly different pay [10, 11] and is maximised by personalisation. The multitask distortion of Section 2.5 arises when high-frequency measurement of instrumentable tasks draws effort from unmeasured ones [7, 8, 9], and real-time systems shorten the horizon over which unmeasured exploratory effort would have to be tolerated. The capacity cost of Section 2.7 arises because imposed income insecurity depletes the very effort the contract was written to elicit [3, 12]. A complete accounting would move the pass-through cases further below the classical baseline than Table 4 shows.

Table 5: What each architecture gets right and wrong

Architecture	Gets right	Gets wrong
Classical periodic pay	Retains a large fixed component, which is correct for a noisy measure; aggregation over the period already averages away much idiosyncratic noise; the pay function is explainable	Leaves the common component in the measure, forgoing the largest available gain; measurement periodicity is inherited from the appraisal calendar rather than chosen
High-frequency measurement without filtering	Reduces idiosyncratic noise; shortens feedback lag; supports faster correction of effort allocation	Runs into a ceiling set by the common share; increases exposure to ratcheting; intensifies the multitask distortion by measuring only what is instrumentable
High-frequency measurement with filtering	Attacks both components; approaches the first best; makes a high variable share genuinely defensible	Requires a defensible peer group, which is contestable; filtering can itself distort effort allocation if applied to variation the agent does influence
Demand-linked pass-through pay	Allocates labour to periods of high demand, which is a genuine allocative function distinct from incentive provision	Transfers common risk to the least able bearer; lowers efficiency below the classical baseline; the variable share moves opposite to the theoretical prescription
Personalised algorithmic pay	Nothing identified in this review on the risk-sharing criterion	Conditions on inferred willingness to accept rather than on contribution; maximises the comparison cost; resists the objective-criteria requirement of pay transparency law
Accelerated payment (earned wage access)	Relieves a genuine within-period liquidity constraint at low cost, and is not credit	Nothing on the risk-sharing criterion, but it is frequently credited with incentive benefits it does not produce

6.4 Timing is not risk sharing: earned wage access

The third dynamic mechanism belongs to a different category and is frequently, and incorrectly, credited to the incentive account. Earned wage access accelerates the payment of wages already earned. It does not alter α, does not alter σ^2, and therefore does not appear anywhere in the expression for η. In the model it is exactly neutral for risk-sharing efficiency.

It is nonetheless welfare-improving, on a margin the classical model does not represent. Constant absolute risk aversion over period income abstracts from liquidity constraints within the period, and it is within-period liquidity that pay timing affects. The regulatory determination that covered products are advances of accrued wages rather than credit [15] is the legal expression of precisely this point: nothing is being lent, and no new risk is being created or transferred. The correct treatment is therefore an additive liquidity term in the joint surplus, independent of α, and the policy implication is a clean separation. A firm may adopt earned wage access on its own merits and should not count it as progress on incentive design; and, symmetrically, an architecture that bundles accelerated payment with demand-linked variable pay is combining a welfare improvement with a welfare loss and should be assessed on the net.

EPRA International Journal of Economic Growth and Environmental Issues- Peer Reviewed Journal
ISSN: 2321-6247
Volume: 14 | Issue: 9 | September 2026 | Journal DOI: 10.36713/epra0713 | SJIF Impact Factor (2026): 8.844

VII. POLICY SYNTHESIS

7.1 For firms designing pay architecture

The design rule that follows from the analysis is a single test applied before adoption: does the proposed architecture raise the ratio of effort-informative variation to total variation in the measure on which pay is conditioned? If it does, the variable share should rise with it, because α^* rises with informativeness. If it does not, the variable share should fall, whatever the sophistication of the system. Firms should compute the common share of variance in their existing performance measures before investing in frequency, because the frequency ceiling of $1/(1 + k_c)$ can be calculated in advance and often makes the investment unattractive on its own terms.

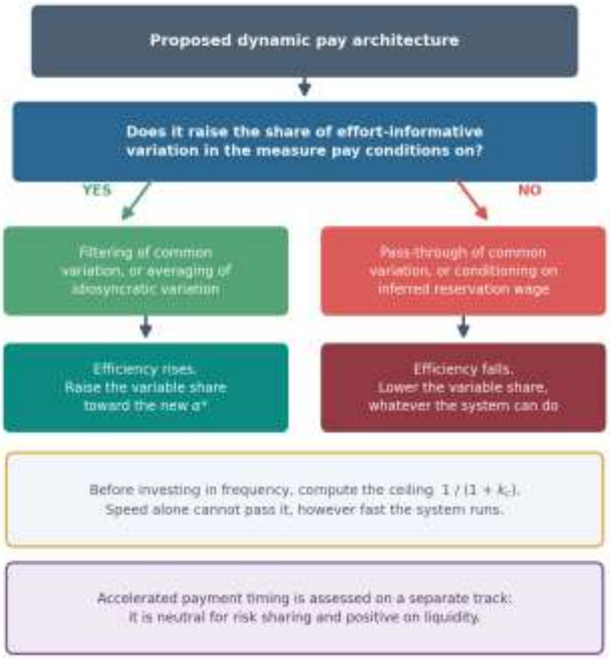

Figure 7: The adoption test — when a dynamic pay architecture improves risk-sharing efficiency

Two cheaper instruments deserve priority over architecture. Filtering within the existing periodic architecture delivers most of the available gain, requires no new technology, and its neglect is the largest single unclaimed efficiency in the field [2]. Widening the controllable share of outcomes through matched delegation of authority and provision of resources [5, 6] raises the efficiency ceiling of every contract the firm can write, and does so without paying a risk premium at all.

7.2 For platform operators and their regulators

Demand-linked real-time pay should be recognised as a risk-transfer instrument rather than an incentive instrument, and assessed accordingly. The distinguishing question is whether the variation passed to the worker is variation the worker can influence: surge pricing that rewards choosing to work at a busy hour is informative about a decision the worker takes, whereas pay that varies with inferred individual willingness to accept is informative about nothing the worker does and is a pure transfer. Kala's evidence that the pay decline was larger for less selective drivers [16] indicates that deployed systems condition on the second. A disclosure requirement covering which variables enter the pay function, without requiring disclosure of coefficients, would separate the two cases at low cost and is a narrower intervention than a general algorithmic audit.

7.3 Governance of learning pay systems

The governance principle appropriate to a pay system that learns is the one Sastry and Arunachalam [17] articulate for learning systems generally: specify the admissible state and action space in advance and require the optimiser to operate inside it, accepting a lower attainable optimum in exchange for excluding states from which recovery is slow. In pay setting this means an ex ante constraint on what the pay function may condition upon — permitting measured contribution and worker-controllable choices, excluding inferred reservation wage, protected characteristics and their proxies — rather than an ex post review of outcomes, which arrives after the dispersion has been created and is in any case difficult to interpret without the counterfactual.

Table 6: An ex ante constraint set for a learning pay system

Constraint	Rationale	Test before deployment
The pay function may condition only on measured contribution and on choices the worker controls	The informativeness principle: variation the worker does not cause carries no information about effort and imposes uncompensated risk	List every input variable and state, for each, the worker action it is informative about
Inferred reservation wage, acceptance propensity and their proxies are excluded	Conditioning on willingness to accept is a transfer, not an incentive; it is the mechanism documented in deployed systems [10, 16]	Confirm that no input is a function of past acceptance or rejection behaviour
Common variation is filtered before it enters pay, not passed through	Filtering is the single largest available efficiency gain; pass-through moves efficiency below the classical baseline	Estimate the common share of variance in the measure and report the frequency ceiling 1/(1 + k_c)
The variable share is set from the informativeness of the measure, not from the capability of the system	The optimal share equals the efficiency ratio; a less informative measure requires a lower, not a higher, variable share	Show that the proposed variable share does not exceed the computed alpha* for the measure used
The pay function is explainable to the worker and to a regulator	Pay transparency law requires differentials for equal work to be justified by objective criteria [18]	Produce a written justification for a representative pay difference between two workers on equal work
Standards are not revised on the same cadence as measurement	Continuous measurement makes ratcheting cheap, and the agent's optimal response is to restrain output [24]	Commit the revision schedule in advance and separate it from the measurement schedule

Regulation is converging on a compatible position from the transparency side. The European Union's Pay Transparency Directive requires pay structures to rest on objective criteria, requires reporting of gaps in variable pay components specifically, and triggers a joint pay assessment where an unexplained gap of five per cent or more persists [18]. An architecture that cannot state why two workers doing equal work were paid differently cannot discharge that obligation. Explainability of the pay function is therefore not an optional ethical enhancement but a condition of lawful operation in a large market, and firms deploying learning pay systems should treat the constraint set as a design input rather than a compliance afterthought.

7.4 For the research community

The decisive unmeasured quantity in this literature is the common share of variance in the performance measures on which pay is actually conditioned. It is estimable from existing personnel and payroll data by decomposing measured performance into a peer-group component and a residual, it determines the frequency ceiling directly, and it is not reported by any compensation survey the review identified. Beyond it, three designs would advance the field: field experiments randomising the filtering of common variation while holding variable share fixed, which isolates the informativeness margin; estimation of the comparison cost of visible pay dispersion for identical work, which the model omits and which personalised architectures maximise; and replication of Kala's staggered design [16] on non-platform employers adopting real-time pay, to establish whether the dispersion result is a property of the architecture or of the platform employment relationship. Table 7 sets out the measurement programme.

Table 7: A measurement programme

What to measure	Why it matters	Feasible source
The common share of variance in the performance measures on which pay is actually conditioned	It determines the frequency ceiling 1/(1 + k_c) directly, and is reported by no compensation survey identified in this review	Firm personnel and payroll records, decomposing measured performance into a peer-group component and a residual
Risk-sharing efficiency computed as the realised variable share against the model-implied alpha*	The gap between the two is the prescription reversal, measured at firm level	Payroll data combined with the variance decomposition above
The comparison cost of visible pay dispersion for equal work	It is omitted from the model and is maximised by personalised architectures	Field or vignette experiments varying the visibility of dispersion at constant expected pay
The effect of filtering at constant variable share	Isolates the informativeness margin from the pay-mix margin, which observational data cannot separate	Field experiment randomising relative performance adjustment across business units
Replication of the staggered algorithmic-pay design outside platform work	Establishes whether the dispersion result is a property of the architecture or of the platform relationship [16]	Employers adopting real-time pay at different dates across comparable sites

VIII. LIMITATIONS AND DIRECTIONS FOR FUTURE RESEARCH

Four limitations qualify these findings. The illustration inherits the assumptions of the linear-exponential-normal framework, and while those assumptions are the ones under which the classical prescriptions were derived, so that the comparison is internally consistent, the absolute efficiency levels reported should not be read as measurements of any real relationship; the ordering of architectures is far more robust than the levels. The decomposition of noise into common and idiosyncratic components treats the split as a known parameter, whereas in practice it must be estimated, and the estimate depends on a peer group whose construction is contestable. The three omitted costs of Section 6.3 are argued to run in one direction but are not quantified, so the pass-through disadvantage is bounded below rather than measured. And the causal evidence on deployed dynamic pay derives principally from one industry and one platform [16], which is the strongest such evidence available but is not a basis for general claims about employers.

Four lines of further work follow. The estimation of the common variance share, set out in Section 7.4, is the priority because it converts the frequency-ceiling result from a comparative statement into an operational one. Second, extending the model to incorporate the comparison cost as a term in the agent's objective would allow personalisation to be priced rather than merely flagged. Third, the interaction between contract design and job design identified in Section 2.4 deserves formal treatment, since the controllable share is a choice variable and no current model treats it as one. Fourth, the ratchet problem is aggravated by continuous measurement in ways this article notes but does not model: a principal who observes performance weekly can revise standards weekly, and the agent's optimal response is to restrain output. Matejka, Mahlendorf and Schäffer [24] provide theory and evidence on the conditions under which target ratcheting induces such restraint, and incorporating it would further reduce the attainable efficiency of high-frequency architectures, since the informational advantage that makes frequent measurement attractive to the principal is the same advantage that makes commitment harder to sustain.

IX. CONCLUSION

Risk-sharing efficiency in the optimal linear contract equals the optimal incentive coefficient, $1/(1 + rc\sigma^2)$, which makes the welfare statistic and the design recommendation the same number and reduces the comparison of pay architectures to a single quantity. Decomposing the noise into common and idiosyncratic components shows that the three mechanisms bundled together under the heading of real-time pay act on different terms and in different directions.

Raising measurement frequency attenuates only idiosyncratic noise and therefore runs into a ceiling set by the common share: in the base case, infinite frequency without filtering cannot exceed an efficiency of 0.455, while filtering at monthly frequency reaches 0.938. Frequency is nearly irrelevant and filtering is nearly everything, and the filtering instrument has been available since the informativeness principle was stated. Passing common variation through to the worker, which is what demand-linked algorithmic pay does, moves efficiency below the unfiltered classical baseline, and the measured behaviour of deployed systems is consistent with that direction. Accelerated payment is welfare-improving on a liquidity margin and exactly neutral on the incentive margin, and should not be credited to the latter.

The practical conclusion is a single test that can be applied before adoption rather than a verdict on technology. A dynamic pay architecture improves risk sharing if and only if it raises the share of effort-informative variation in the measure on which pay is conditioned. Nothing about a system's speed, granularity or sophistication reveals which side of that line it falls on, and the observed prescription reversal — the highest variable shares attached to the least filtered measures — indicates that the line is not currently being drawn. The governance implication follows directly: what a pay algorithm may condition upon has to be settled in advance, because by the time the dispersion is visible in the outcome data, the surplus has already been destroyed.

REFERENCES

1. *Holmstrom, B. (2016). Pay for Performance and Beyond. Prize Lecture in Economic Sciences in Memory of Alfred Nobel, The Royal Swedish Academy of Sciences, Stockholm.*
2. *Bertrand, M., & Mullainathan, S. (2001). Are CEOs rewarded for luck? The ones without principals are. Quarterly Journal of Economics, 116(3), 901-932.*
3. *Cameron, L. D., & Meuris, J. (2026). The perils of pay variability: Determinants of worker aversion to variable compensation in low- and middle-wage jobs. PLOS ONE, 21(5), e0347194.*
4. *Sastry, N. S. K., & Selvarasu, A. (2014). A study of impact on performance appraisal on employee's engagement. International Journal of Managerial Studies and Research, 2(11), 10-22.*
5. *Sastry, N. S. K., & Ravish, P. Y. (2025). Empowering employees in the Indian workplace: The transformative power of delegation and resources. EPRA International Journal of Research and Development, 10(7).*
6. *Sastry, N. S. K., & Ravish, P. Y. (2025). The impact of delegated authority and resource provision on decision-making: Fostering employee initiative. EPRA International Journal of Multidisciplinary Research, 11(7).*

7. *Holmstrom, B., & Milgrom, P. (1991). Multitask principal-agent analyses: Incentive contracts, asset ownership, and job design. Journal of Law, Economics, and Organization, 7(Special Issue), 24-52.*
8. *Sastry, N. S. K. (2024). An impact of departmental structure on task achievement. EPRA International Journal of Environmental, Economics, Commerce and Educational Management, 11(9).*
9. *Sastry, N. S. K., & Mallya, M. M. (2025). Economic constraints and the status quo: Barriers to innovation and creative problem-solving. EPRA International Journal of Research and Development, 10(9), 191-197.*
10. *Dubal, V. (2023). On algorithmic wage discrimination. Columbia Law Review, 123(7), 1929-1992.*
11. *Sastry, N. S. K. (2025). Talent designation and social comparison: Economic drivers of organizational growth - an empirical analysis. EPRA International Journal of Economics, Business and Management Studies, 12(10).*
12. *Sastry, N. S. K. (2024). Revitalizing leadership to combat workforce burnout. EPRA International Journal of Environmental, Economics, Commerce and Educational Management, 11(9).*
13. *Sastry, N. S. K., & Mallya, M. M. (2025). Revolutionizing HR hiring with AI analytics: Enhancing efficiency and attracting top talent. EPRA International Journal of Research and Development, 10(9), 125-131.*
14. *Fortune Business Insights. (2026). Earned Wage Access Market Size, Share and Industry Analysis, 2026-2034. Fortune Business Insights, Pune. [Commercial market research; figures used as orders of magnitude.]*
15. *Consumer Financial Protection Bureau. (2025). Truth in Lending (Regulation Z); Non-application to Earned Wage Access Products. Advisory Opinion, 90 Fed. Reg., effective 23 December 2025.*
16. *Kala, J. (2025). Algorithmic Wage Discrimination: Evidence from Uber. Working paper, Department of Economics, University of California, Riverside.*
17. *Sastry, N. S. K., & Arunachalam, T. (2025). Probabilistic safe RL: Modelling and managing risk in dynamic environments. EPRA International Journal of Research and Development, 10(9), 83-87.*
18. *European Parliament and Council. (2023). Directive (EU) 2023/970 to strengthen the application of the principle of equal pay for equal work or work of equal value through pay transparency and enforcement mechanisms. Official Journal of the European Union; transposition deadline 7 June 2026.*
19. *Page, M. J., McKenzie, J. E., Bossuyt, P. M., Boutron, I., Hoffmann, T. C., Mulrow, C. D., and others. (2021). The PRISMA 2020 statement: An updated guideline for reporting systematic reviews. BMJ, 372, n71.*
20. *Pay Governance LLC. (2026). S&P 500 CEO Compensation Trends. Harvard Law School Forum on Corporate Governance, Cambridge, MA.*
21. *Holmstrom, B., & Milgrom, P. (1987). Aggregation and linearity in the provision of intertemporal incentives. Econometrica, 55(2), 303-328.*
22. *Prendergast, C. (1999). The provision of incentives in firms. Journal of Economic Literature, 37(1), 7-63.*
23. *Lazear, E. P. (2000). Performance pay and productivity. American Economic Review, 90(5), 1346-1361.*
24. *Matejka, M., Mahlendorf, M. D., & Schaffer, U. (2024). The ratchet effect: Theory and empirical evidence. Management Science, 70(1), 128-142.*