Supplementary Information

Metabolic labeling of RNA uncovers principles of RNA production and degradation dynamics in mammalian cells

Michal Rabani, Joshua Z Levin, Lin Fan, Xian Adiconis, Raktima Raychowdhury, Manuel Garber, Andreas Gnirke, Chad Nusbaum, Nir Hacohen, Nir Friedman, Ido Amit and Aviv Regev

Supplementary Notes

1. Measuring RNA transcription rate by 4-thiouridine (4sU) metabolic labeling of newly-transcribed RNA in mouse DCs

We isolated newly-transcribed as well as total cellular RNA from mouse DCs that were metabolically labeled with 4sU, and quantified it using both qPCR and nCounter. We added 4sU to DCs at specific time-points following stimulation with LPS, for a pre-defined labeling time (**Supplementary Fig. 1a**); RNA molecules that were actively transcribed during that time are labeled by 4sU (**Supplementary Fig. 1b**). We isolated the entire cellular RNA population (**RNA-Total**), used an *in vitro* reducing chemical reaction to specifically and covalently link biotin to 4sU residues, and separated the 4sU labeled RNA (**RNA-4sU**) using biotin capture with streptavidin magnetic beads (**Supplementary Fig. 1c**). Finally, we quantified the RNA levels in both populations using either standard qRT-PCR for 5 genes or direct capture and count of individual transcripts by the nCounter system for a signature of 254 genes (**Supplementary Table 1; Methods**) that are key representatives of the LPS response based on our previous study (Amit et. al., Science, 2009).

Both the biotinylation and 4sU incorporation have an efficiency of less than 100%. Since in our computational models we base our estimates on relative RNA measurements and absolute RNA concentrations, our transcription rates may possibly be under-estimates.

1.1. 4sU metabolic labeling is specific, reproducible, consistently measured by qRT-PCR and nCounter, and has no significant effect on cellular function or transcriptional response of primary DCs.

We assessed the performance of 4sU labeling by several measures. **First**, without 4sU labeling, only minimal background RNA levels are evident following 4sU purification, whereas after 1 hour labeling, LPS-induced genes are on average 6.8-fold enriched in labeled vs. total RNA, while un-induced but constantly expressed controls are only 2.5-fold enriched on average, as expected for stable genes. Longer labeling time (4h labeling) increased the average fold enrichment of LPS-induced genes to 10.8-fold (**Supplementary Fig. 2a,b**). Thus, our method specifically labels newly transcribed RNA. Furthermore, after a sufficiently long time, all the RNA in the cell will be labeled, and labeled RNA will be equal to total RNA (ratio of ~1). Although 24h labeling is not enough to reach that point, we see that after such long labeling the differences between induced and control genes decrease, since a higher fraction of the RNA population reaches its turnover time. The nCounter measurements are consistent with qRT-PCR across LPS-induced genes (**Supplementary Fig. 2c,d**). Fold enrichment versus the flow-through fraction is even higher than versus total RNA, indicating robust and specific capture (**Supplementary Fig. 3**). **Second**, measurements from replicate experiments are highly correlated, demonstrating reproducibility of the labeling and purification protocols (**Supplementary Fig. 4**). **Third**, total RNA levels do not significantly change following labeling of up to 24 hours, suggesting minimal effect of 4sU on transcriptional responses in DCs, as reported for other cell types (Dolken et. al. RNA, 2008) (**Supplementary Fig. 5**).

1.2. After 10 min labeling RNA-4sU is predominantly nuclear.

We extracted RNA-4sU separately from nuclear and cytoplasmic fractions, and quantified each fraction. On average, more than 70% of mRNA-4sU is present in the nuclear fraction at short labeling times (10 min) and only ~50% in longer labeling times (45 min) (**Supplementary Fig. 6a**).

1.3. Promoter binding by RNA polymerase II (Pol-II) peaks at or before RNA-4sU measurements following short labeling.

We used Pol-II ChIP to measure binding of Pol-II to promoters of 12 selected genes. Although the binding of Pol-II at promoters does not directly estimate transcription rates, changes in binding may still indirectly reflect transcriptional activity. We find that promoter binding by RNA polymerase II (Pol-II) peaks at or before RNA-4sU measurements following short labeling (**Fig. 1b**, **Supplementary Fig. 6b**), supporting the short-4sU metabolic labeling as a direct measurement of RNA transcription rates.

2. Testing the fit of the 'constant degradation' and 'varying degradation' models to the nCounter data

We use the 'goodness of fit' test (**Methods**) to measure how well the data fits the predictions of a given model. Rejecting the model in this case means that one of the modeling assumptions does not hold in the data, but we cannot determine which assumption that is. The 'constant degradation' model fits the majority of genes well, and we reject the 'constant degradation' model ($P < 0.01$) in 16% of the cases (42/254) (**Supplementary Fig. 8a**). The 'varying degradation' model improves the fit to the data compared with the 'constant degradation' model, and it is rejected ($P < 0.01$) in only 8% of the cases (21/254) (**Supplementary Fig. 8b**).

However, the 'varying degradation' is more complex than the 'constant degradation' model. Therefore, when fitting the 'varying degradation' model to the data we search through a much wider range of possible degradation functions, but this range includes within it also the constant functions. Consequently, if indeed a (relatively) constant function provides a good fit for the data, the 'varying degradation' model can still predict that. Indeed, for most genes that accept the 'constant degradation' model, the 'varying degradation' model fits them with a largely constant function, such that the results of both models are very similar on this set of genes, with a somewhat better fit by the 'varying degradation' model.

Although the varying degradation model fits the data slightly better (more genes are rejected by the 'constant degradation' model than the 'varying degradation' model), this marginal improvement must be weighted versus the preference for a simpler model. The

'varying degradation' model allows for many degrees of freedom (12 parameters in two 'impulse model' functions are fit from 26 data points), which may lead to over-fitting and may be highly sensitive to noise. Conversely, the 'constant degradation' model only fits 7 free parameters.

To better compare the models, we therefore used the 'likelihood ratio test' (**Methods**) to decide when the improvement is big enough to justify the increase in complexity of the model (measured as the number of parameters). In most cases, the simpler model ('constant degradation') is favored, but we find 44 genes, which significantly ($p<0.01$) reject the 'constant degradation' model in favor of the 'varying degradation' model. Notably, the varying degradation rate profiles calculated for these genes are smooth and involve a few (one or two) prominent changes in degradation rates (**Fig. 2d,e**), suggesting that they indeed reflect reliable estimates.

Finally, we note that the distribution of p-values from the likelihood ratio test is highly bi-modal, and our results are very robust to the specific p-value threshold chosen. For example, taking $p < 0.05$ rather than $p < 0.01$ has a minimal effect on the result, increasing the number of genes that reject the 'constant degradation' model only by 19 genes (an additional 8%), and resulting in 63 (rather than 44) of the 'signature set' genes rejecting the model.

3. The degradation rates predicted by the model are supported by using Actinomycin D

Using transcriptional arrest by Actinomycin D (actD) is the prevalent method in the literature to measure RNA degradation rates. Although actD experiments have inherent limitations, in particular for measuring degradation rates in a dynamic setting, such as during LPS response, we used this standard protocol to compare our results to the accepted gold-standard method used today.

We treated cells with actD either before (0h) or at 2.5 hours after LPS stimulation (**Methods; Supplementary Fig. 11-12**) and measured mRNA levels at 6 or 4 time points following the treatment (respectively). We confirmed actD activity and transcriptional arrest with a mock experiment that measured RNA levels at the same time points, in cells which were not treated with actD (data not shown).

Although the actD method is commonly used to measure RNA degradation rates, several lines of evidence show that this method is significantly biased and unreliable, particularly in dynamic settings. First, the transcriptional arrest induced by actD is an enormous stress on living cells, and therefore RNA levels following such treatment are significantly biased by it (Pelechano and Perez-Ortin., Yeast, 2008). Second, two basic assumptions of the standard analysis (see **Methods**) are problematic: (1) studies show that actD blocks only about 95% of the transcriptional activity (Shutt and Krueger, J Immunol, 1972), and (2) it is also unclear whether or not the degradation rate remains constant following actD addition.

Our approach, combining metabolic labeling and modeling, avoids many of the biases of using actD. It does not stress the cell and does not interfere with the cellular machinery or with the normal dynamics of the response, however it infers the RNA degradation rates indirectly.

We compared our approach to the standard actD experiments in two different ways.

First, we compared actD measurements to the mRNA levels predicted by either the 'constant' or 'varying' degradation models under the same conditions (i.e. when assuming no new transcription from the respective time point onward, and using the degradation rates estimated by each model from our analysis). At 0h, for the 83% of genes that accept the 'constant degradation' hypothesis, the predictions of both models fit the actD measurements to the same extent, but for the remaining 17%, the 'varying degradation' model fits the actD data substantially better, supporting our conclusions (**Supplementary Fig. 11**). At 2.5h, the differences between the models are minimal for all genes, since the predicted constant rates are dominated by the (later) higher rates and therefore both models predict very similar rates at this time point.

Second, we used the actD measurements to predict the degradation rates at 0h and at 2.5h (separately, see **Methods**), and compared these to the rates predicted by the models. This analysis leads to similar conclusions, showing that the genes that reject the 'constant degradation' hypothesis correlate better with the 'varying degradation' model predictions, while for the accepting genes the correlation with the 'constant degradation' model is higher (**Supplementary Fig. 12a,b**). Furthermore, the genes that reject the 'constant degradation' model also showed bigger differences in degradation rates estimated based

on each of the two actD experiments (**Methods; Supplementary Fig. 12c**), typically increasing at 2.5h.

Thus, the actD experiments, despite all their biases, support our findings.

4. Comparing library composition between 4sU-Seq and standard RNA-Seq (polyA+, RNA-Total)

In the absence of a preliminary polyA selection step, the 4sU-Seq libraries measure a significantly higher level of rRNA. Most of these reads are not uniquely mapped to the genome (since rRNA genes are repetitive), and therefore we mapped them independently directly to rRNA sequences (**Methods**). Of the 4sU-Seq library reads, ~30% map to rRNA sequences (vs. 0.4% in RNA-A+-Seq) and only ~60% uniquely map to the genome (vs. ~82% in RNA-A+-Seq) (**Supplementary Tables 3-4**). Although without polyA selection rRNA levels rise as expected, they still remain significantly lower than their estimated proportion of the overall cellular RNA (~90%).

More generally, the 4sU-Seq libraries measure a substantially different RNA population than standard (polyA+, RNA-Total) RNA-Seq libraries (**Fig. 3b; Supplementary Fig. 15**). Mappable reads are distributed differently across the genome. While the RNA-A+-Seq libraries are mainly enriched for mRNA exons (80% vs. 18%), the 4sU-Seq libraries are enriched with other RNA populations, including rRNA (14.2% vs. 0.2%), introns of pre-mRNAs (47% vs. 8.1%) and pre-miRNAs (0.1% vs. 0.02%). Among the genomic features with the most substantial bias in favor of RNA-A+-Seq (**Supplementary Fig. 16**), are annotated exons (77%) and other non-annotated regions (14.5%), which generate polyadenylated transcripts (either protein coding or non-coding RNAs). Conversely, bias in favor of 4sU-Seq is very common in introns (87%), rRNA (~1%) and non-annotated regions (10.5%). This latter population may reflect transient, unannotated, non-polyadenylated transcripts, such as antisense or non-coding RNAs.

Differences between the libraries arise both since 4sU-Seq measures newly-transcribed RNA, and since it avoids polyA selection biases (such as increased 3' read coverage, and elimination of RNA species without polyA signals). For example, the replication-dependent Histone genes, which are the only eukaryotic mRNAs that lack a polyA tail (Marzluff, Curr Opin Cell Biol, 2005) are not detectable in the RNA-A+-Seq samples but are clearly and differentially expressed in 4sU-Seq (**Supplementary Fig. 17a**). Additionally, several primary microRNA transcripts (pri-miRNAs) are substantially more abundant in 4sU-Seq than in RNA-A+-Seq (**Supplementary Fig. 17b**). Although pri-miRNA are poly-adenylated (Cai et. al. RNA, 2004), efficient pri-miRNA processing occurs co-transcriptionally, before poly-adenylation (Pawlicki and Steitz, Trends Cell Biol, 2010). Thus, some pri-miRNA transcripts may be short-lived, consistent with their enrichment in 4sU-Seq.

5. Comparing different RNA degradation rate estimates

We use several comparisons to verify the consistency of the degradation rates we estimated, both within our own data and with other methods or other works in different systems.

First, for the 254 'signature set' genes, we compared the two independent data sources we obtained: the nCounter high-resolution data and the genome-wide 4sU-Seq and RNA-A+-Seq measurements. We find that the constant degradation rates predicted by our model in each of the two data sets (**Figure 3c** and **Supplementary Figure 10a**) are highly similar (data not shown). Moreover, of the 44 genes with predicted varying degradation based on the signature-set data, 28 (64%) also reject the 'constant degradation' model based on the genome-wide data (although the statistical tests we use are different), while remaining 16 genes peaks are not detectable at the 1h time resolution of the genomic experiment.

Second, for the 254 'signature set' genes, we compared the degradation rates estimated by the constant and varying degradation models with the rates we predict based on the Actinomycin D data (**Methods**; **Supplementary Figure 12a-b**). We find that the two measurements are highly correlated, but that our models predict *shorter* RNA half-lives on average (i.e., higher degradation rate) than half-lives measured by Actinomycin D (e.g., median half-life at 0 hours is 26 min. with the 'varying degradation' model, and 80 min. with Actinomycin D). The inherited biases in standard Actinomycin D techniques may contribute to the observed discrepancies. Our measurements indicate that Actinomycin D treatment does not block transcription completely (data not shown), and

that can leads to under-estimated degradation rates (i.e., predicts the RNA is more stable), as indeed we find when comparing to our model predictions.

Finally, we compared our model's genome-wide predictions to RNA degradation rates measured previously in fibroblasts by 1 hour of metabolic labeling (Dolken et. al. RNA, 2008). We again find that the ranking of estimated half-lives in our experiments positively correlates (Spearman ρ=0.56) with fibroblasts degradation rates (**Supplementary Fig. 20**), but our half-life estimates are *shorter* on average. Some of the discrepancies here could be accounted for by the improved RNA quantification assays we use compared with Dolken et. al., and by differences in cell type. Indeed, there are evidence in the literature (Leclerc et. al., Cancer Cell Int., 2002) for twice as long RNA half-life in one cell type versus another.

Notably, these discrepancy in absolute degradation rates may be due to inherent limitations in all techniques (ours, which estimates degradation rates indirectly, as well as others, which are significantly biased), and highlight the need for direct and reliable methods to measure degradation rates. Yet, despite the possible discrepancy in the absolute degradation rates, the correlation with other rate estimates remains high, suggesting that the ranking of the RNAs stability and most importantly, the ratio between RNA transcription and degradation rates, on which we base our conclusions, are reliable and accurate.

6. Alternative isoforms and overlapping transcripts

Our model and analysis are conducted entirely at the level of a gene locus. Thus, if multiple splice isoforms are expressed simultaneously, our computed transcription, processing and degradation rates represent an 'average' over all the transcripts from that locus.

To assess whether the signature genes and other transcripts in our system express more than one major RNA species during the response to LPS, we applied Scripture, an *ab-initio* transcript reconstruction algorithm developed in our lab (Guttman et. al. Nature Biotechnology. 2010) to our strand-specific polyA+ RNA-Seq libraries. We considered all splicing isoforms with minimal expression (based on RPKM) and a significant difference (more than 300 bases) between them.

In the **signature set**, we find that 24% (60/254) of our signature set genes express only one major RNA species throughout the response. Furthermore, examining the probes designed for genes with more than one splice variant, we find that for 70% of genes (178/254), the probe would detect over 90% (i.e., missing at most 1 variant) of RNA species produced from that gene (including both pre-mRNA and mRNA). This is consistent with the intended design of the code set, and indicates that it can capture the needed RNA species for the vast majority (94%, 238/254) of the genes.

In the **genome-wide data**, approximately 27% (2,737/10,106) of the Refseq genes express only one major transcript, and over 75% express less than 5 (between 1-4) transcripts. We also find evidence for antisense transcription in 3% (268/10,106) of the RefSeq genes.

Although many of the genes in our system express alternative isoforms, our model does not distinguish overlapping transcripts, as noted above. Thus, if several alternative RNA species are produced from a single gene, each with a different degradation rate and a different temporal expression pattern (i.e., different RNA variants dominate the RNA population at different times), then the predicted variable degradation rates might actually reflect that change. However, other factors can also lead to variation between specific sub-populations of RNAs, including cellular localization (e.g. nuclear vs. cytoplasmic, specific locations in the cytoplasm, or RNA sequestration in p-bodies) and time of production (e.g. very 'old' RNAs are marked by proteins and stabilized). Since distinguishing sub-populations within a specific gene's RNA population is beyond the scope of our current model, we consider the degradation rate that we predict as an 'average' degradation rates for the entire population. If the relative frequency of stable splice variants will change over time, so will the 'average' degradation rate of the population, which is exactly what our model predicts.

Alternative isoforms can also lead to biases in estimating pre-mRNA levels. Since we base our estimates on distinguishing exonic from intronic reads, if an exon in one isoform is excised with an intron in another, in a manner not reflected in the canonical (Refseq) definition of the gene's exons and introns, then our estimates can be inaccurate. Moreover, since our libraries are not strand-specific, transcription from the other strand (most likely antisense-transcription) can also lead to similar biases. We therefore performed several tests and introduced several filters to address these potential issues.

1. We included in this analysis **only transcripts with a single splice variant in Refseq.** However, Refseq likely underestimates the number of isoforms (see e.g. Trapnell et

al, Nature Biotechnology 2010), and our RNA-Seq data likely reflects additional splice variants and transcripts.

2. **Alternative isoforms do not have a major impact on our RPKM estimates.** We examined each exon defined by our *ab-initio* predicted transcripts. An 'exon' is defined as any transcript that is expressed in the polyA+ libraries (our transcript assembler ensures these are flanked by appropriate splice donor and acceptor sites, and supported by 'spliced' reads). We then asked how many introns in the Refseq annotations overlap such an exon (over at least 50% of the intron's length) in either the sense or antisense strand. We expect that biases caused by overlapping transcription will lead to different RPKM estimates between introns of the same transcript. We therefore compared the RPKM (from 4sU-Seq) of adjacent introns in the same transcript that either (**1**) do not intersect any exons (either sense or antisense), (**2**) only one of which intersects an exon, or (**3**) both intersect an exon. A one-sided test suggests that the values in group (2) or (3) are not significantly higher than in group (1) ($p < 0.98$ and $p<0.62$ respectively), thus supporting our estimated pre-mRNA RPKM values.
3. **Extreme alternative isoforms biases are limited.** Although alternative isoforms biases do not seem to have a global effect, they might still have a major impact in specific cases. We therefore looked for expressed transcripts that overlap a sense or antisense strand exon in 10% or more of their introns. We find that out of the expressed transcripts, 16% (1469/10,106) overlap a sense exon, and 1% (84/10,106) overlap an antisense exon, suggesting that such extreme cases are limited to a small fraction of the transcripts.

In light of these findings, **we excluded from our analysis transcripts with significant alternative isoforms biases.** When fitting the 'constant degradation' model, we removed from the set of 10,106 expressed RefSeq genes all genes for which we find evidence for antisense transcription (268 genes; 3%). When fitting the 'constant degradation and processing' model, we excluded from the set of 3,011 expressed genes (with sufficient intron and exon expression) all genes with antisense transcription (125 genes; 4%) and all genes that overlap a sense or an antisense exon (764 genes; 25%). We conservatively analyzed only the remaining 2,122 genes, for which our predictions should be relatively accurate. Notably, genes with multiple isoforms are enriched with ribosomal proteins, and genes with anti-sense transcription are enriched with immune response genes (chemokines), including many TFs (**Supplementary Figure 24c**).

Supplementary Methods

nCounter code sets design

Code sets were designed and constructed to detect the 254 'signature' genes. In designing the gene specific probes, NanoString utilizes the RefSeq database as the source for transcript sequences. The target sequence of interest is scanned from 5' to 3' for all possible 100nt target regions. Each region is screened and scored for the following properties: sequence composition (%GC, polyC, repeat sequences, non-ATGC bases), homology with other targets in the transcriptome, thermodynamic properties (Tm), and isoform coverage (if applicable). The selection algorithm selects a single probe for the target. The location of the probe is not biased by position along the transcript, but rather is directed by the most optimal sequence from a homology and thermodynamic perspective, and if possible – to constitutive exons, capturing multiple splice isoforms.

Each probe matches a 100 bases long exonic sequence of the target genes, and therefore detects both pre-mRNA and mature mRNA. The location of these 100 bases within the exonic sequence of the gene varies uniformly across the 254 genes in our signature set, ranging from probes located at the very beginning (5' end) of the gene, and some located at the 3' end of the gene.

nCounter data normalization

We normalized the nCounter data (**Supplementary Table 2**) in two steps. **First**, we controlled for small variations in the efficiency of the automated sample processing by normalizing all samples analyzed on a given run to the levels of the sum of positive spiked-in controls provided by the nCounter system. **Second**, we relied on 8 control

genes (Ppp2r1a, Ndufs5, Psma7, Tomm7, Psmb4, Ndufa7, Eif4h, Capza1), which are highly expressed and were identified from previous work[1] as unperturbed upon LPS stimulation. For every sample, we computed the weighted average m_i of the RNA counts of the seven transcripts, and normalized the sample's values by multiplying by $1/m_i$.

Estimating the 'fraction of explained variance'

To estimate the percent of variance in a set of measurements $\{y_1 ... y_n\}$ that is explained by a model predictions $\{x_1 ... x_n\}$, we use the following formulation.

First, define the following quantities (where $\bar{y} = \frac{1}{n}\sum_{i=1}^{n} y_i$ is the mean of the measurements):

$SS_{Total} = \sum_{i=1}^{n}(y_i - \bar{y})^2 =$ the total sum of squares

$SS_{\text{Regression}} = \sum_{i=1}^{n}(x_i - \bar{y})^2 =$ the explained sum of squares

$SS_{Error} = \sum_{i=1}^{n}(y_i - x_i)^2 =$ the residual sum of squares

If for our regression model $SS_{Total} = SS_{Error} + SS_{\text{Regression}}$ (e.g. linear regression), than the unexplained variance is:

$$\text{fraction of explained variance} = \frac{SS_{\text{Regression}}}{SS_{Total}} = \frac{SS_{Total} - SS_{Error}}{SS_{Total}} = 1 - \frac{SS_{Error}}{SS_{Total}} = r^2$$

where r^2 is the 'coefficient of determination' of the regression model.

However, this is not true in our case. Therefore, we define the difference by $SS_U = SS_{Total} - (SS_{Error} + SS_{\text{Regression}})$, and now we get that $SS_{Total} = SS_{Error} + SS_{\text{Regression}} + SS_U$, and so:

$$\text{fraction of explained variance} \leq \frac{SS_{\text{Regression}}}{SS_{Total} + |SS_U|}$$

Our estimator is therefore an upper bound on the fraction of explained variance.

Dynamic model of RNA transcription and degradation in the signature set

Modeling. We use a first-degree dynamic model that directly connects RNA expression with transcription and degradation rates. Formally, let α be the transcription rate (RNA/min*cell) and β the degradation rate (1/min) and let *X* be the expression level of a gene *x* (RNA/cell), the time evolution of *X* in time is directly determined by the rates:

$$\frac{dX}{dt} = \alpha(t) - \beta(t)X$$

When we measure RNA-total we globally integrate the rates over the entire lifetime of the cell. Therefore:

$$X(0) = X_0$$
$$X(T) = \int_{-\infty}^{T} \alpha(t) - \beta(t)X\,dt = X_0 + \int_{0}^{T} \alpha(t) - \beta(t)X\,dt$$

However, when we measure $\tilde{X}$, the RNA-4sU level of gene *x* after labeling for t_L minutes, we locally integrate the rates:

$$\tilde{X}(T - t_L) = 0$$
$$\tilde{X}(T) = \tilde{X}_0 + \int_{0}^{T} \alpha(t) - \beta(t)\tilde{X}\,dt = \int_{T-t_L}^{T} \alpha(t) - \beta(t)\tilde{X}\,dt$$

To estimate α, we make two assumptions: **(1)** α and β are approximately constant during the (short) labeling period, and **(2)** with short enough labeling time ($\leq 10\,\text{min}$), RNA-4sU is mostly nuclear, and hence we assume it is subjected to little if any degradation (β=0).

With the first assumption we have an analytic solution for the differential equation (where α_T and β_T are the constant values for α and β at labeling time T):

$$\tilde{X}(T) = \frac{\alpha_T}{\beta_T}\left(1 - e^{-\beta_T \cdot t_L}\right)$$

and when β approaches 0 (the second assumption) we get that:

$$\tilde{X}(T) = \frac{\alpha_T}{\beta_T}\left(1 - e^{-\beta_T \cdot t_L}\right) \overset{\beta_T \to 0}{\to} t_L \alpha_T$$

$$\Rightarrow \alpha_T = \frac{\tilde{X}(T)}{t_L}$$

Therefore, $\tilde{X}$ directly measures the average transcription rate during the labeling period.

We use a modeling scheme that compares two parametric descriptions of the rate functions. Since both the transcription rate (α, RNA-4sU) and the expression levels (X, RNA-total) are measured experimentally, we expect both to include a certain level of noise. We assume an additive Gaussian noise model, which we estimated from experimental repeats. Fitting smooth functions (like a constant or a sigmoid function) to the data, should remove much of this noise, and yield more accurate results.

We describe temporally varying rates with the 'impulse' model, a 6-parameter double-sigmoid function[2,3]:

$$rate(t) = \frac{1}{h_1}\left(h_0 + \left(h_1 - h_0\right)\frac{1}{1 + e^{-\lambda(t - t_1)}}\right)\left(h_2 + \left(h_1 - h_2\right)\frac{1}{1 + e^{-\lambda(t - t_2)}}\right)$$

The first model ('**constant degradation**') uses an impulse model for the transcription rate *α(t)* and a simple one-parameter constant function for the degradation rate. The second model ('**varying degradation**') uses a different parameterization and estimates a temporally varying $\beta(t)$, which is modeled using a second impulse model.

Fitting the model to the data. The likelihood of the data $\{y_1 ... y_n\}$, given the model predictions $\{x_1 ... x_n\}$ is:

$$\log L(D;M) = \log L(\{y_1 ... y_n\};\{x_1 ... x_n\},\sigma_{Exp},\sigma_{Tx}) = \log \prod_{i=1}^{n} p(y_i \mid x_i,\sigma) = \sum_{i=1}^{n} \log p(y_i \mid x_i,\sigma)$$

With each model, we optimize this target function using the Nelder-Mead simplex algorithm, a popular direct search method for multidimensional unconstrained minimization that does not use analytic gradients[4].

Our optimization procedure includes several steps:

(1) We fit an impulse model to α(t) (optimized by multiple random initializations):

$$\alpha(t) = \frac{\tilde{X}(t)}{t_L} = \frac{\tilde{X}(t)}{10}$$

(2) We fit a second impulse model directly to RNA-total data (X(t)).

(3) We use the functions fitted to α(t) and X(t) to find an initial estimate of β(t). We use the analytic solution and relating X(t) and α(t) to β(t)

$$\beta(t) = \frac{1}{X}\left(\alpha(t) - \frac{dX}{dt}\right)$$

and estimate β(t) at measured time points. We treat these estimated values as traning data, and find the parameters (either for constant function or impulse model) that optimize the fit to these values.

(4) Using the initial parameters of α(t) and β(t) as initialization points, and jointly optimize the parameters to maximize the likelihood of the model given the entire data (including both α(t) and X(t)).

Statistical tests. We use two complementary statistical tests. The first is a **goodness of fit test** that measures how well the data fits each model (separately). Rejecting the model in this case means that one of the modeling assumptions does not hold in the data, but we cannot determine which assumption that is. We estimate the variance from data points measured in replicates, and use the standard least square error and associated chi-square distribution with *n* degrees of freedom[5] to calculate a p-value:

$$Var(y|x) = \sum_{i=1}^{n} \frac{(x_i - y_i)^2}{\sigma^2} \Rightarrow$$
$$pvalue = P\left(y' \middle| Var(y' \mid x) > Var(y \mid x)\right) \sim \chi^2(n)$$

The second test is a **likelihood ratio test**, which compares the two alternative models, and identifies in which cases the simpler model ('constant degradation') should be rejected in favor of the more complex model ('varying degradation'). We use a standard nested likelihood ratio test[6].

Read mapping

All reads were aligned to the mouse reference genome (NCBI 37, MM9) using the TopHat aligner[7] with default parameters. Briefly, TopHat uses a two-step mapping process, first using Bowtie[8] to align all reads that map directly to the genome (with no gaps), and then mapping all reads that were not aligned in the first step using gapped alignment. TopHat initially splits each read into 25 bases long segments, and align each segment independently (with up to 2 mismatches), and in the next step merge these segments together into the complete read alignment. To identify split reads (splice junctions) it requires an 'anchor region' of 8 bases mapped with no mismatches on each

side of the gap. We used only reads that were mapped uniquely to the genome for further analysis.

Quantification of transcript abundance from RNA-Seq data

We estimate the expression of a transcript *X* in both RNA-A+-Seq and RNA-4sU-Seq by standard 'Reads Per Kilobase exon model per Million mapped reads' (RPKM), as previously described in[9], but we define it over exons alone:

$$RPKM(X) = \frac{10^9 \cdot r(X_{Exon})}{R_{Exon} \cdot s(X_{Exon})}$$

where $r(X_{Exon})$ is the number of reads mapped to exons of *X*, R_{Exon} is the sum of reads mapped to all exons (of all transcripts) in the experiment, and $s(X_{Exon})$ is the total length (in kilobases) of exons in X.

Since the RNA-A+-Seq and 4sU-Seq libraries measure markedly different RNA populations (**Fig. 3b**), by concentrating on exonic reads alone we are able to compare the expression of genes in both types of libraries.

We normalized RPKM estimates in two steps. **First**, we select a set of control genes (NM_010885, NM_018753, NM_013477, NM_011116, NM_011873, NM_008138, NM_007505, NM_010684) as previously described[10]. For every sample, we computed the weighted average m_i of the RPKM of the controls, and normalized the sample's values by multiplying by $1/m_i$. **Second**, the mean and standard deviation of each sample's log(*RPKM*) values is adjusted to the overall mean and standard deviation across all samples. We used measurement of the ratio between quantity of RNA in each sample before and after labeled RNA purification (0.046 ± 0.019 SD), to normalize RNA-total and RNA-4sU samples relative to each other.

We identify expressed genes by taking genes with $\log_2$(RPKM) >= *2* for at least a single time point. Of the 25,000 mRNAs in the Refseq[11] database, ~40% (10,106) are expressed above threshold for at least one time point in both RNA-4sU and RNA-A+-Total. From this set of genes, we excluded transcripts with significant antisense expression (to avoid biases arising from strand-specificity; see **Supplementary Note 6**), leaving us with 9,838 genes.

Dynamic modeling of RNA transcription and degradation in genome-wide data

Compared to our signature-set measurements, in our genome-wide data we expanded the temporal scope and reduced the temporal resolution, generating a 6h time course, with 1h intervals between measurement points. To extract labeled RNA in sufficient quantities for sequencing, we had to extend the labeling time (to 45 min). As labeling time increases RNA-4sU is increasingly subjected to degradation, and thus no longer directly approximates transcription rates, but a 'local integration' of the average transcription and degradation rates over 45 minutes. All these changes affect our modeling scheme, and we adjust it accordingly.

First, the lower temporal resolution, and the fewer time points do not allow fitting a model with too many parameters. Therefore, we cannot fit the 'varying degradation' model in this setting, and only fit the 'constant degradation' model (that has 7 parameters). As a result, we can only use the 'goodness of fit' test (above), and find genes that reject the 'constant degradation' model, but we cannot determine which of the model's assumptions lead to the rejection.

Second, we use our signature-set predictions to guide the search for optimal parameters for the genomic data, by directing the search towards parameters used in the signature set. To this end, we build a Gaussian mixture model prior on the model parameters using the optimal signature-set parameters. We estimate one Gaussian from each of the clusters I-VIII, resulting with a mixture of 8 Gaussian models for each parameter. We adapt the likelihood function to include the prior:

$$\log L(D;M) = \log p(M)p(D \mid M) = \log p(M) + \log p(D \mid M)$$

Where the second term is estimated as before, and the first term is the prior:

$$P(M) = p(\{\theta_j\}) = \frac{1}{8}\sum_{i=1}^{8} P(\{\theta_j\} \mid \mu_i, \sigma_i)$$

$$P(\{\theta_j\} \mid \mu_i, \sigma_i) = \prod_{j=1}^{m} N(\theta_j \mid \mu_{ij}, \sigma_{ij})$$

We optimize the likelihood function as before.

Finally, since we no longer have replicated measurements, we estimate the variance of the data based on the background distribution of expression values in the genome, and use this estimate in the 'goodness of fit' test.

Quantification of mRNA and pre-mRNA abundance from 4sU-Seq data

We estimate the expression of the pre-mRNA of *X* by intron-RPKM, and the overall (pre-mRNA+mRNA) expression of *X* by exon-RPKM:

$$intronRPKM(X) = \frac{10^9 \cdot r(X_{Intron})}{(R_{Exon} + R_{Intron}) \cdot s(X_{Intron})}$$

$$exonRPKM(X) = \frac{10^9 \cdot r(X_{Exon})}{(R_{Exon} + R_{Intron}) \cdot s(X_{Exon})}$$

where *r(X)* is the number of reads mapped to exon or intron sequences of *X*, *R* is the sum of reads mapped to all exon or intron sequences of all transcripts in the experiment, and *s(X)* is the total length (in kilobases) of exon or intron sequences in X.

We implicitly assume that intronic reads are predominantly derived from pre-mRNAs, rather than from excised introns. In theory, some of these reads (those that overlap only introns, rather than intron-exon junctions) can arise from sequencing of the excised introns and lead to over-estimated pre-mRNA expression levels and under-estimated processing rates. However, equivalent reads coverage of introns and intron-exon junctions in our 4sU-Seq data suggest that this is unlikely. Furthermore, previous studies also suggest that introns degrade very fast (e.g. 3.5 and 6 min for two example introns[12]). Finally, deletion of the de-branching enzyme (Dbr1 in yeast or its mouse ortholog) lead to fast accumulation of high levels of circular introns[13-15], suggesting that during normal growth de-branching and intron degradation occur very fast.

Although a high fraction of the reads (47%) map to introns, since introns are very long, their coverage is relatively low (**Supplementary Fig. 12b**, the intron annotation normalized both by library size and by length in kb). Therefore, we focus on the 30% of the expressed genes (3,011/10,106) in which both introns and exons are expressed above a threshold (RPKM >= 2) for at least one time point. From this set of genes, we excluded transcripts with antisense transcription or significant alternative isoforms biases (see **Supplementary Note 6**), leaving us with 2,122 genes.

Dynamic modeling of processing rates

We estimated the expression of recently transcribed pre-mRNAs and mature mRNAs separately, by distinguishing exonic from intronic reads (see above). Intronic reads are produced only from pre-mRNA transcripts, and therefore their expression estimates the pre-mRNA-4sU level (P). Conversely, exonic reads are produced from both mature mRNA and pre-mRNA, and therefore their expression measures mRNA-4sU, the entire recently transcribed RNA population (X), which is the sum of pre-mRNA (P) and mature mRNA (M).

Modeling. Let α be the transcription, β the degradation and γ the processing rate. We extend our previous dynamic model to include the processing rate as following:

(1) $\frac{dP}{dt} = \alpha - \gamma P$

(2) $\frac{dM}{dt} = \gamma P - \beta M$

(3) $X = M + P \quad \Rightarrow \frac{dX}{dt} = \alpha - \beta M$

In this model we assume that **(1)** mRNA is produced in two steps: first pre-mRNA is transcribed, and then it is processed into mature mRNA, and **(2)** that only the mature mRNA is exported to the cytoplasm, where it is degraded. Therefore, we describe processing rates in terms of pre-mRNA half-life (while mRNA half-life refers to degradation rate). We use a modeling scheme with an 'impulse' model[2,3] for transcription rate, and with constant degradation and processing rates.

Since we use 4sU-Seq data, we locally integrate the rates:

$$P(T) = \int_{T-t_L}^{T} \alpha(t) - \gamma P \, dt$$

$$M(T) = \int_{T-t_L}^{T} \gamma P - \beta M \, dt$$

To estimate α, we make two assumptions: **(1)** pre-mRNA is at steady state during the (short) labeling period, and **(2)** all rates (α, β and γ) are approximately constant during the (short) labeling period.

From the first assumption, we can simplify equation (2) of the model:

$$\frac{dP}{dt} = \alpha - \gamma P = 0$$
$$\Rightarrow \gamma P = \alpha$$
$$\Rightarrow \frac{dM}{dt} = \gamma P - \beta M = \alpha - \beta M$$

And, from the second assumption, we have an analytic solution for this equation:

$$M(T) = \frac{\alpha_T}{\beta_T}\left(1 - e^{-\beta_T \cdot t_L}\right)$$
$$\Rightarrow \alpha_T = M(T) \cdot \frac{\beta_T}{1 - e^{-\beta_T \cdot t_L}}$$

To estimate α, we use here the temporally constant β values that we estimated with the previous model.

Fitting the model to the data. We optimize the likelihood of the data as described before. We fit the parametric description to the estimated α(t) and P(t) values using equation (1) of the model. Our optimization procedure includes several steps:

(1) We fit an impulse model to α(t) (optimized by multiple random initializations).

(2) We fit a constant to initial values for γ(t). To predict initial γ(t) values, we assume that during the (short) labeling period β=0 and all other rates (α, β) are constant. Therefore:

$$\frac{dX}{dt} = \alpha - \beta M = \alpha \quad \Rightarrow X(T) = \alpha \cdot T$$
$$\frac{dP}{dt} = \alpha - \gamma P \qquad \Rightarrow P(T) = \frac{\alpha}{\gamma}\left(1 - e^{-\gamma T}\right)$$
$$\frac{P(T)}{X(T)} = \frac{1 - e^{-\gamma T}}{\gamma T} = f(\gamma T)$$
$$\Rightarrow \gamma = \frac{1}{T} f^{-1}\left(\frac{P(T)}{X(T)}\right)$$

(3) We use these parametric descriptions of α(t) and γ as initialization points, and optimize equation (1) of the model globally over both P(t) and α(t).

Testing the fit. We use a goodness-of-fit test as before (with the standard least square error). When rejecting the model, the goodness-of-fit test cannot indicate which of the modeling assumption led to rejection. Specifically here, we cannot say if it is the constant degradation or the constant processing rate assumptions (or both) that lead to the rejection.

1 Amit, I. *et al.* Unbiased reconstruction of a mammalian transcriptional network mediating pathogen responses. *Science* **326**, 257-263 (2009).
2 Chechik, G. & Koller, D. Timing of gene expression responses to environmental changes. *J Comput Biol* **16**, 279-290 (2009).
3 Chechik, G. *et al.* Activity motifs reveal principles of timing in transcriptional control of the yeast metabolic network. *Nat Biotechnol* **26**, 1251-1259 (2008).
4 Lagarias, J. C., Reeds, J. A., Wright, M. H. & Wright, P. E. Convergence properties of the Nelder-Mead simplex method in low dimensions. *Siam J Optimiz* **9**, 112-147 (1998).
5 Hodges, J. L. & Lehmann, E. L. *Basic concepts of probability and statistics*. 2d edn, (Holden-Day, 1970).
6 Wilks, S. S. The large sample distribution of the likelihood ratio for testing composite hypotheses. *Annals of Mathematical Statistics*, 60-70 (1938).

7 Trapnell, C., Pachter, L. & Salzberg, S. L. TopHat: discovering splice junctions with RNA-Seq. *Bioinformatics* **25**, 1105-1111 (2009).
8 Langmead, B., Trapnell, C., Pop, M. & Salzberg, S. L. Ultrafast and memory-efficient alignment of short DNA sequences to the human genome. *Genome Biol* **10**, R25 (2009).
9 Mortazavi, A., Williams, B. A., McCue, K., Schaeffer, L. & Wold, B. Mapping and quantifying mammalian transcriptomes by RNA-Seq. *Nat Methods* **5**, 621-628 (2008).
10 Vandesompele, J. *et al.* Accurate normalization of real-time quantitative RT-PCR data by geometric averaging of multiple internal control genes. *Genome Biol* **3**, RESEARCH0034 (2002).
11 Pruitt, K. D., Tatusova, T. & Maglott, D. R. NCBI reference sequences (RefSeq): a curated non-redundant sequence database of genomes, transcripts and proteins. *Nucleic Acids Res* **35**, D61-65 (2007).
12 Clement, J. Q., Qian, L., Kaplinsky, N. & Wilkinson, M. F. The stability and fate of a spliced intron from vertebrate cells. *RNA* **5**, 206-220 (1999).
13 Nam, K., Lee, G., Trambley, J., Devine, S. E. & Boeke, J. D. Severe growth defect in a Schizosaccharomyces pombe mutant defective in intron lariat degradation. *Mol Cell Biol* **17**, 809-818 (1997).
14 Kim, Y. J., Bjorklund, S., Li, Y., Sayre, M. H. & Kornberg, R. D. A multiprotein mediator of transcriptional activation and its interaction with the C-terminal repeat domain of RNA polymerase II. *Cell* **77**, 599-608 (1994).
15 Kim, H. C., Kim, G. M., Yang, J. M. & Ki, J. W. Cloning, expression, and complementation test of the RNA lariat debranching enzyme cDNA from mouse. *Mol Cells* **11**, 198-203 (2001).

Supplementary Figure 1

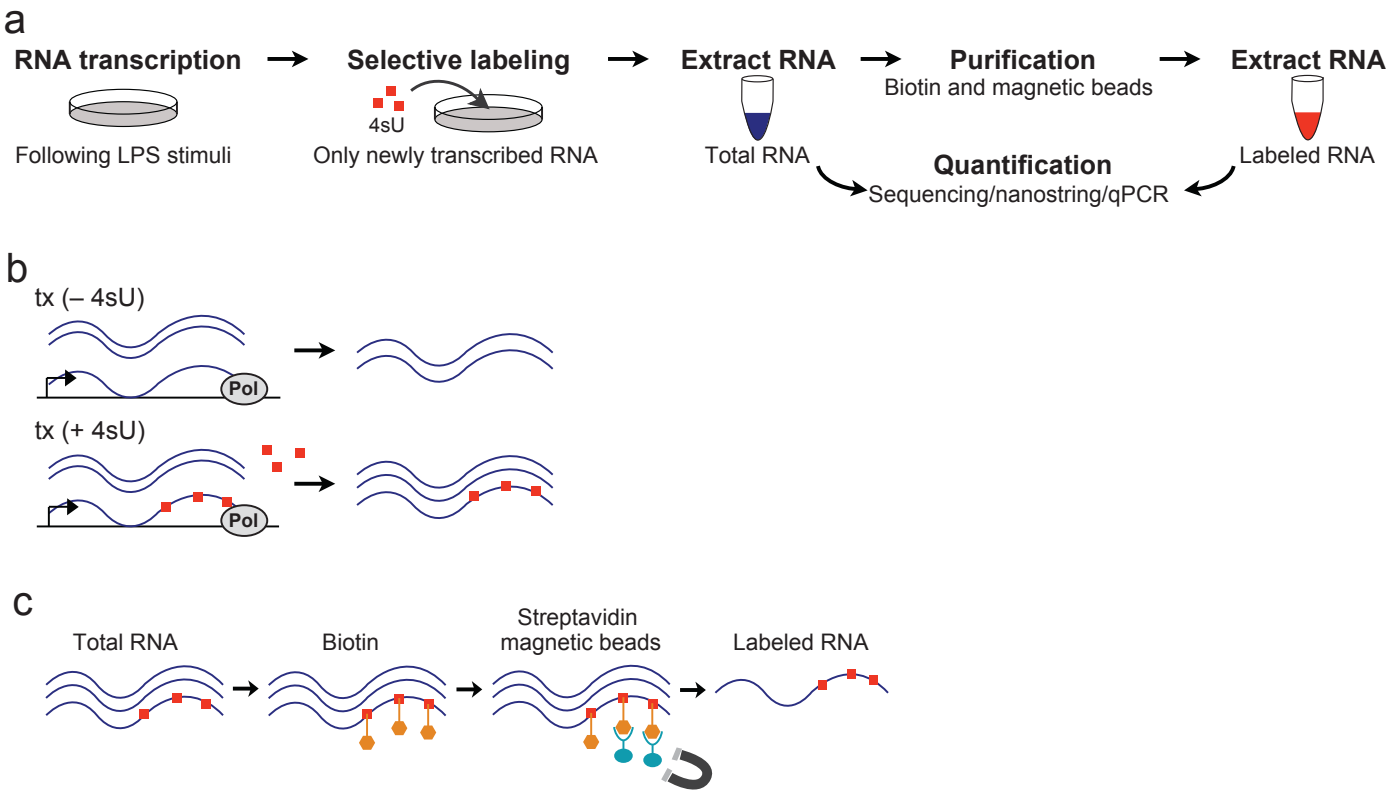

.

Supplementary Figure 1: Metabolic labeling of newly transcribed RNA by 4-thiouridine (4sU).

(a) Metabolic labeling using 4-thiouridine (4sU). After stimulation, we add 4sU (red squares) to growing cells for a pre-defined time, collect the cells and extract total RNA (blue). Using biotin capture with streptavidin magnetic beads, we purify labeled RNA (red) from the total RNA extract. **(b)** Transcription with and without 4sU. When 4sU is present, it is incorporated into the growing RNA chain in place of uridine. **(c)** Purification using streptavidin magnetic beads. Total RNA extract is biotinylated by covalently linking biotin (2, orange) to 4sU, followed by binding to Streptavidin coated magnetic beads (3, light blue). Biotylinated (4sU labeled) RNA is magnetically isolated, whereas unlabeled RNA is washed out. Finally, cleaving the biotin-4sU disulfide bond releases the labeled RNA from the beads (4).

Supplementary Figure 2

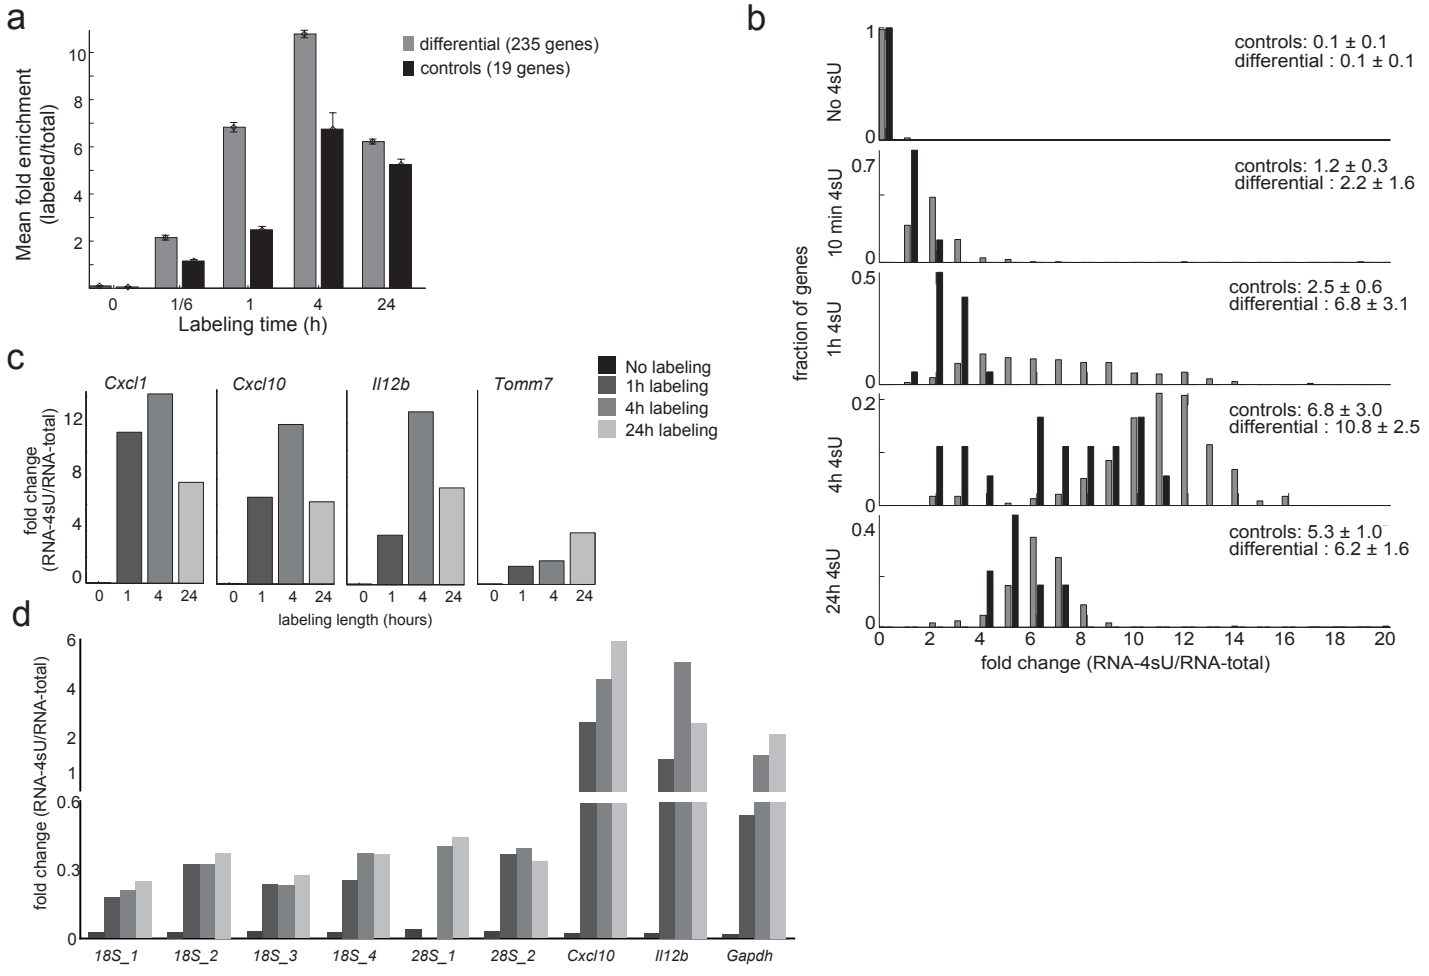

.

Supplementary Figure 2: Newly transcribed (4sU labeled) RNA is enriched in responsive genes in DCs at 4h after LPS stimulation.

Shown is the fold enrichment in 4sU labeled vs. total RNA in DCs, at 4h post-LPS stimulation, after different metabolic labeling times (0h: no labeling; 10 min and 1, 4 or 24 hours of metabolic labeling). **(a)** Newly transcribed (4sU labeled) RNA is enriched in responsive genes in DCs at 4h after LPS stimulation. Shown are the mean fold enrichments in 4sU labeled vs. total RNA (Y axis) based on nCounter quantification. Black: 19 control genes whose expression does not change during the response to LPS; gray: 235 genes whose expression changes in the response (either induced or repressed). Error bars represent the standard error. **(b)** nCounter quantification of 19 control genes (black), whose expression does not change during the response to LPS and of 235 genes (gray), whose expression changes in the response (either induced or repressed). Shown is the distribution of fold enrichment in 4sU labeled vs. total RNA. The mean and standard deviation of the fold change for control and induced genes is indicated to the right. **(c)** Four illustrative examples of fold enrichment in 4sU labeled vs. total RNA (Y-axis, measured by nCounter) after different labeling times (X-axis). Gene names shown on top (Three induced genes: Cxcl10, Il12b, Cxcl1, and one control gene: Tomm7). **(d)** Four illustrative examples of fold enrichment in 4sU labeled vs. total RNA (Y-axis, measured by qRT-PCR) after different labeling times (colored bars, ordered from no labeling to 24h). Gene names shown on bottom (two induced genes: Il12b, Cxcl10, one control gene: Gapdh and two rRNA controls: 18S, 28S).

Supplementary Figure 3

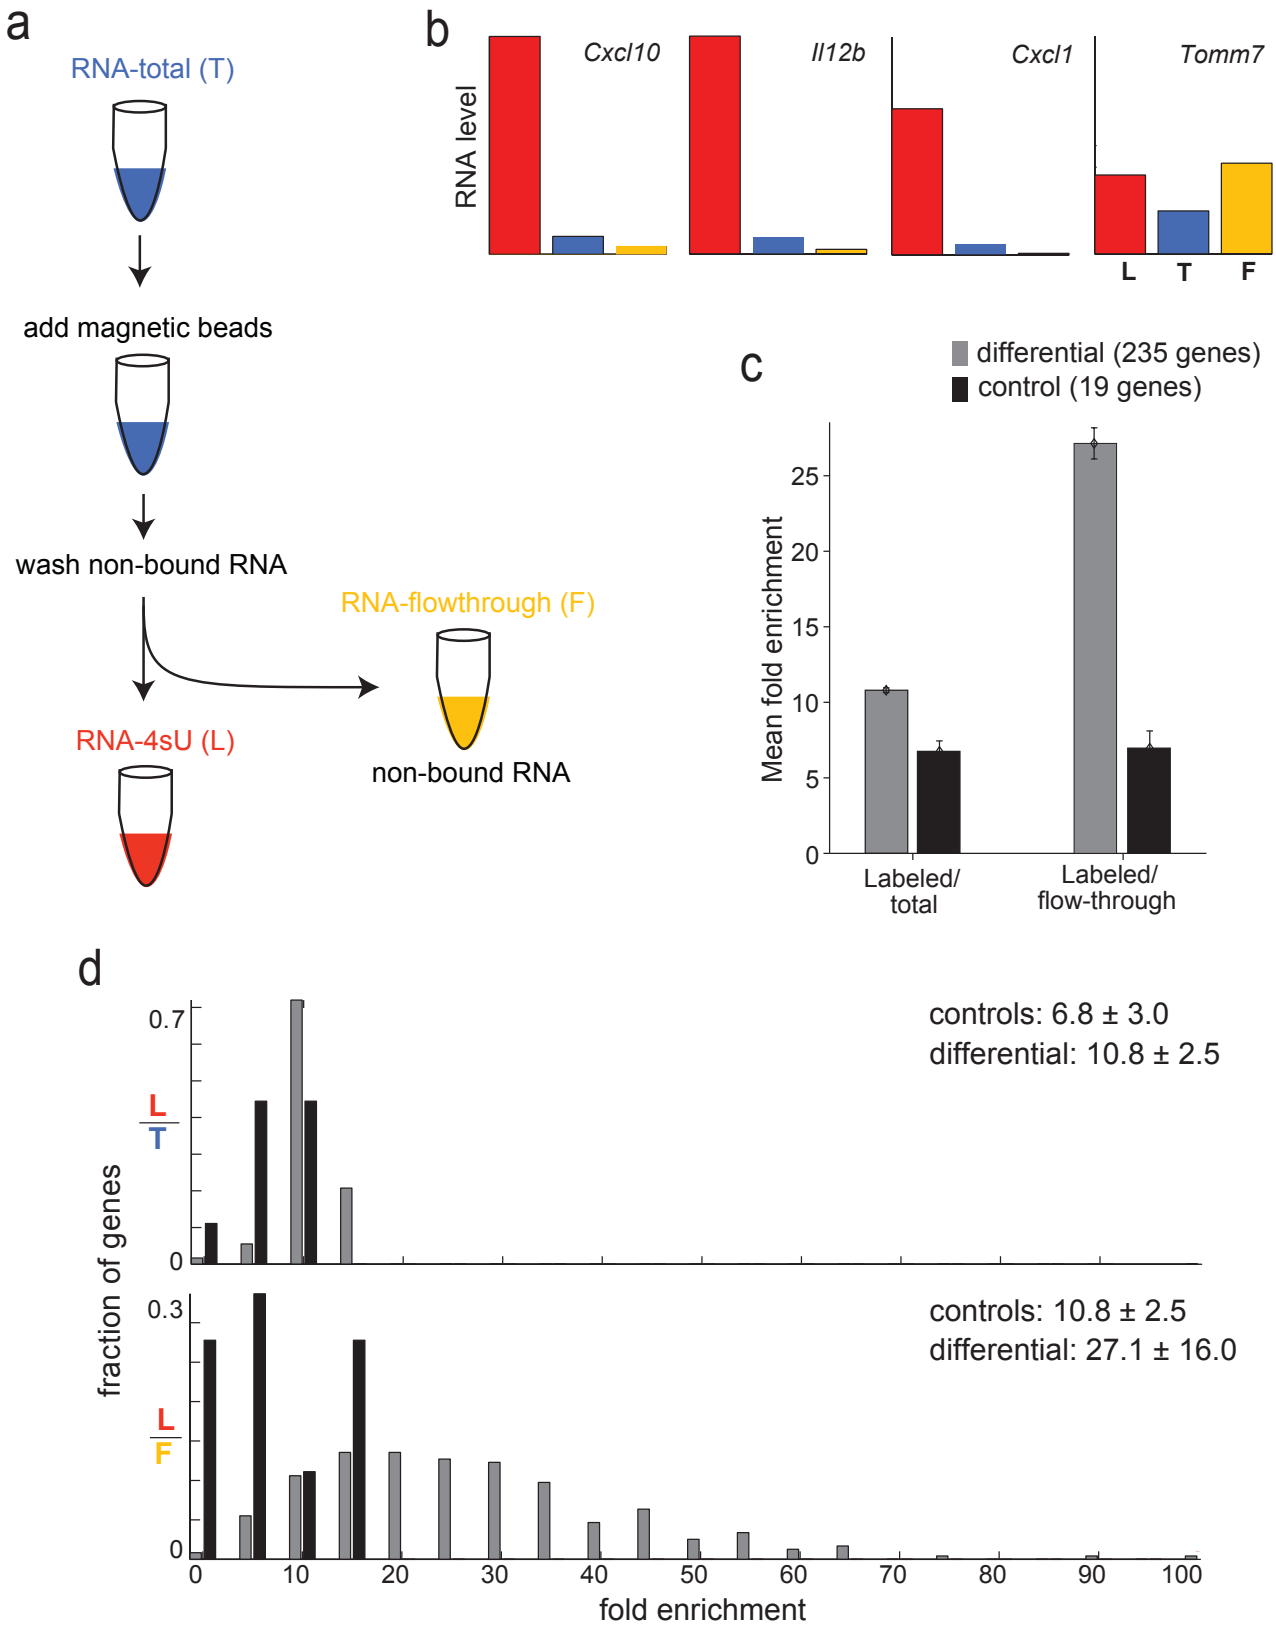

Supplementary Figure 3: 4sU labeled RNA is enriched relative to flow through RNA.

(a) Experiment overview. We add 4sU to growing cells for a pre-defined labeling time, collect the cells and extract total RNA (T, blue). Using biotin capture with streptavidin magnetic beads, we purify labeled RNA (L, red) from the total RNA extract. Labeled RNA (L) binds to the magnetic beads, while the rest of the RNA is washed, and collected separately (flow-through RNA, F, yellow). All data collected at 4 hours after LPS stimulation and 4 hours 4sU labeling. **(b)** Expression level of four example genes (3 induced genes: Cxcl10, Il12b, Cxcl1, and one control gene: Tomm7) as measured by nCounter from total RNA (T, blue), 4sU labeled RNA (L, red) and flow-through RNA (F, yellow). **(c)** 4sU purified RNA is enriched relative to flow-through RNA. Shown is the mean fold enrichment (Y axis) of 4sU-purified RNA vs. total RNA (left) or vs. flow-through RNA (right). Black: 19 control genes, whose expression does not change during the response to LPS; Grey: 236 genes, whose expression changes in the response (either induced or repressed). Error bars represent the standard error. **(d)** Distribution of fold enrichment of RNA-4sU over RNA-Total (top) or flow-through RNA (bottom), based on nCounter quantification. The mean and standard deviation fold change for all control and induced genes is indicated to the right.

Supplementary Figure 4

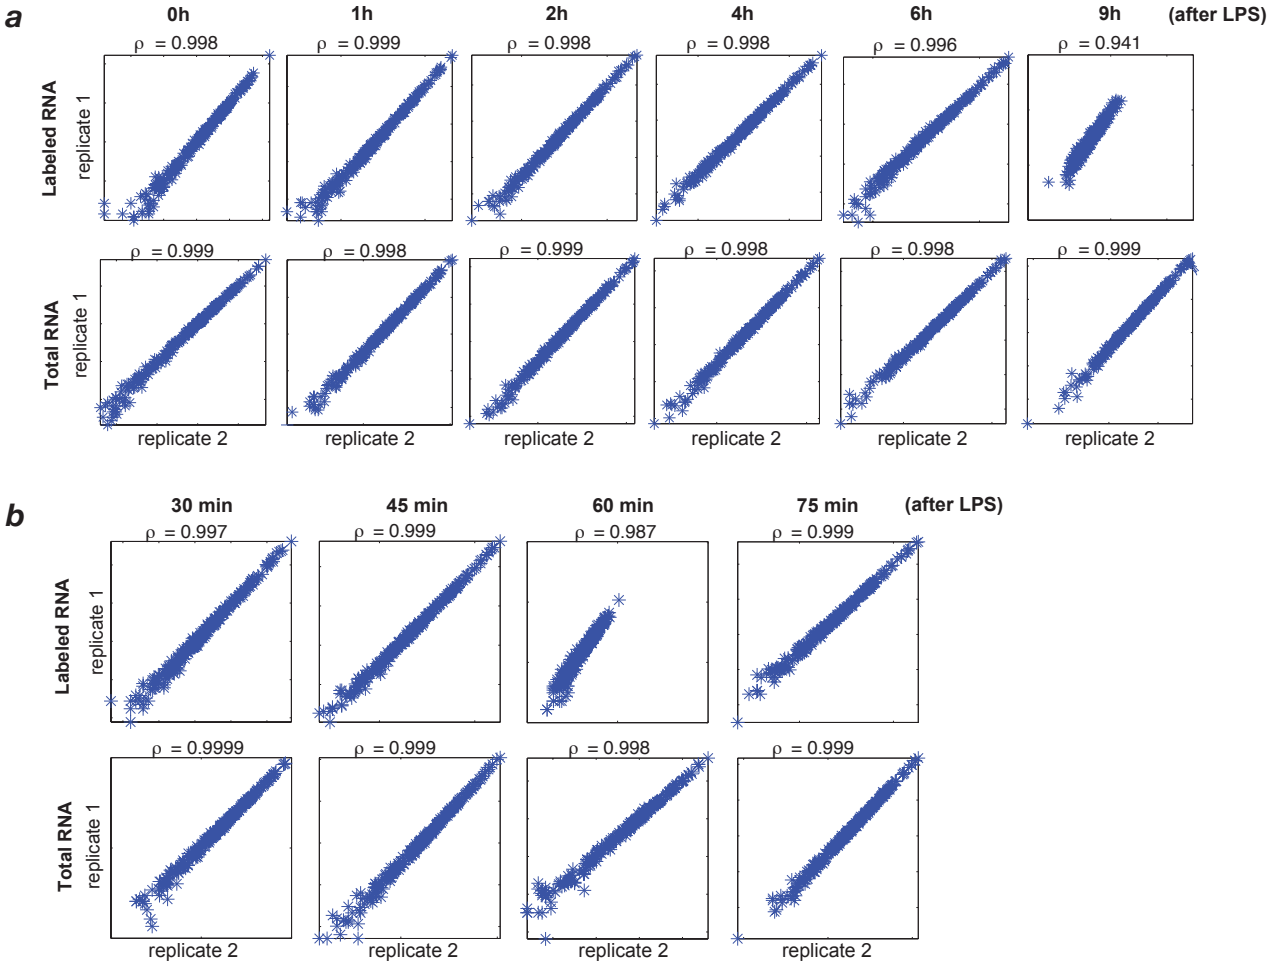

.

Supplementary Figure 4: nCounter measurements of RNA-4sU and RNA-Total are reproducible.

Shown are scatter plots of nCounter measured expression values (X and Y axes, log scale) for 254 signature genes from replicate experiments, for either RNA-4sU (top row) or RNA-Total (bottom row), along with the Pearson correlation coefficient (ρ, top). Replicates represent two samples per time point for which all the purification steps were conducted independently per sample. **(a)** RNA-4sU labeled and RNA-Total collected in DCs at different times post-LPS stimulation (0h: no stimulation; 1, 2, 4, 6 or 9 hours after stimulation) after 30 min metabolic labeling. **(b)** RNA-4sU and RNA-Total collected in DCs at different times post-LPS stimulation (30, 45, 60 or 75 minutes after stimulation), after 10 min metabolic labeling.

Supplementary Figure 5

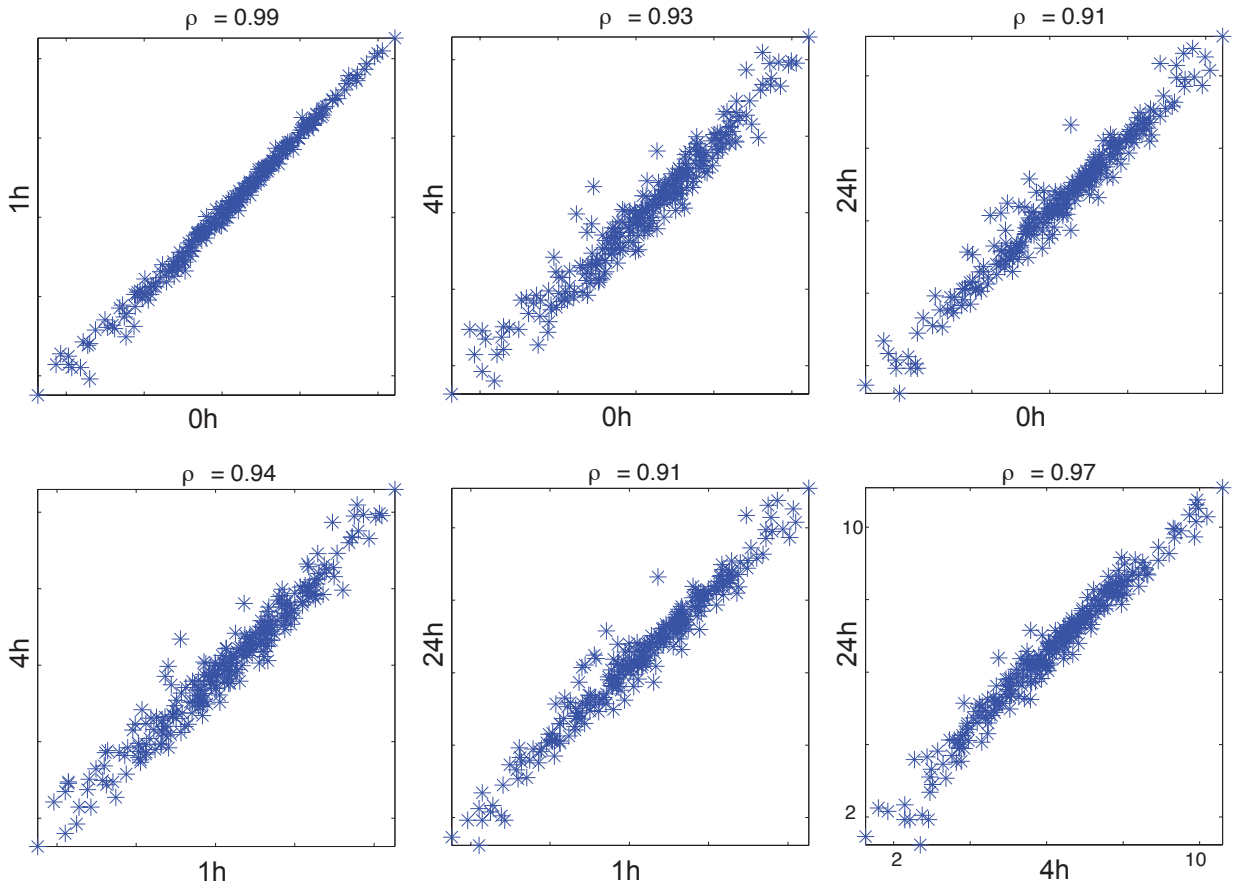

.

Supplementary Figure 5: Minimal effect of growth in 4sU medium on expression levels in response to LPS stimulation.

Shown are scatter plots of RNA-Total expression values (X and Y axes, log scale) measured by nCounter for 254 signature genes, at 4h post-LPS stimulation, after different metabolic labeling times (0h: no labeling; 1, 4 or 24 hours of metabolic labeling, axis labels). The Pearson correlation coefficient (ρ) is shown on top.

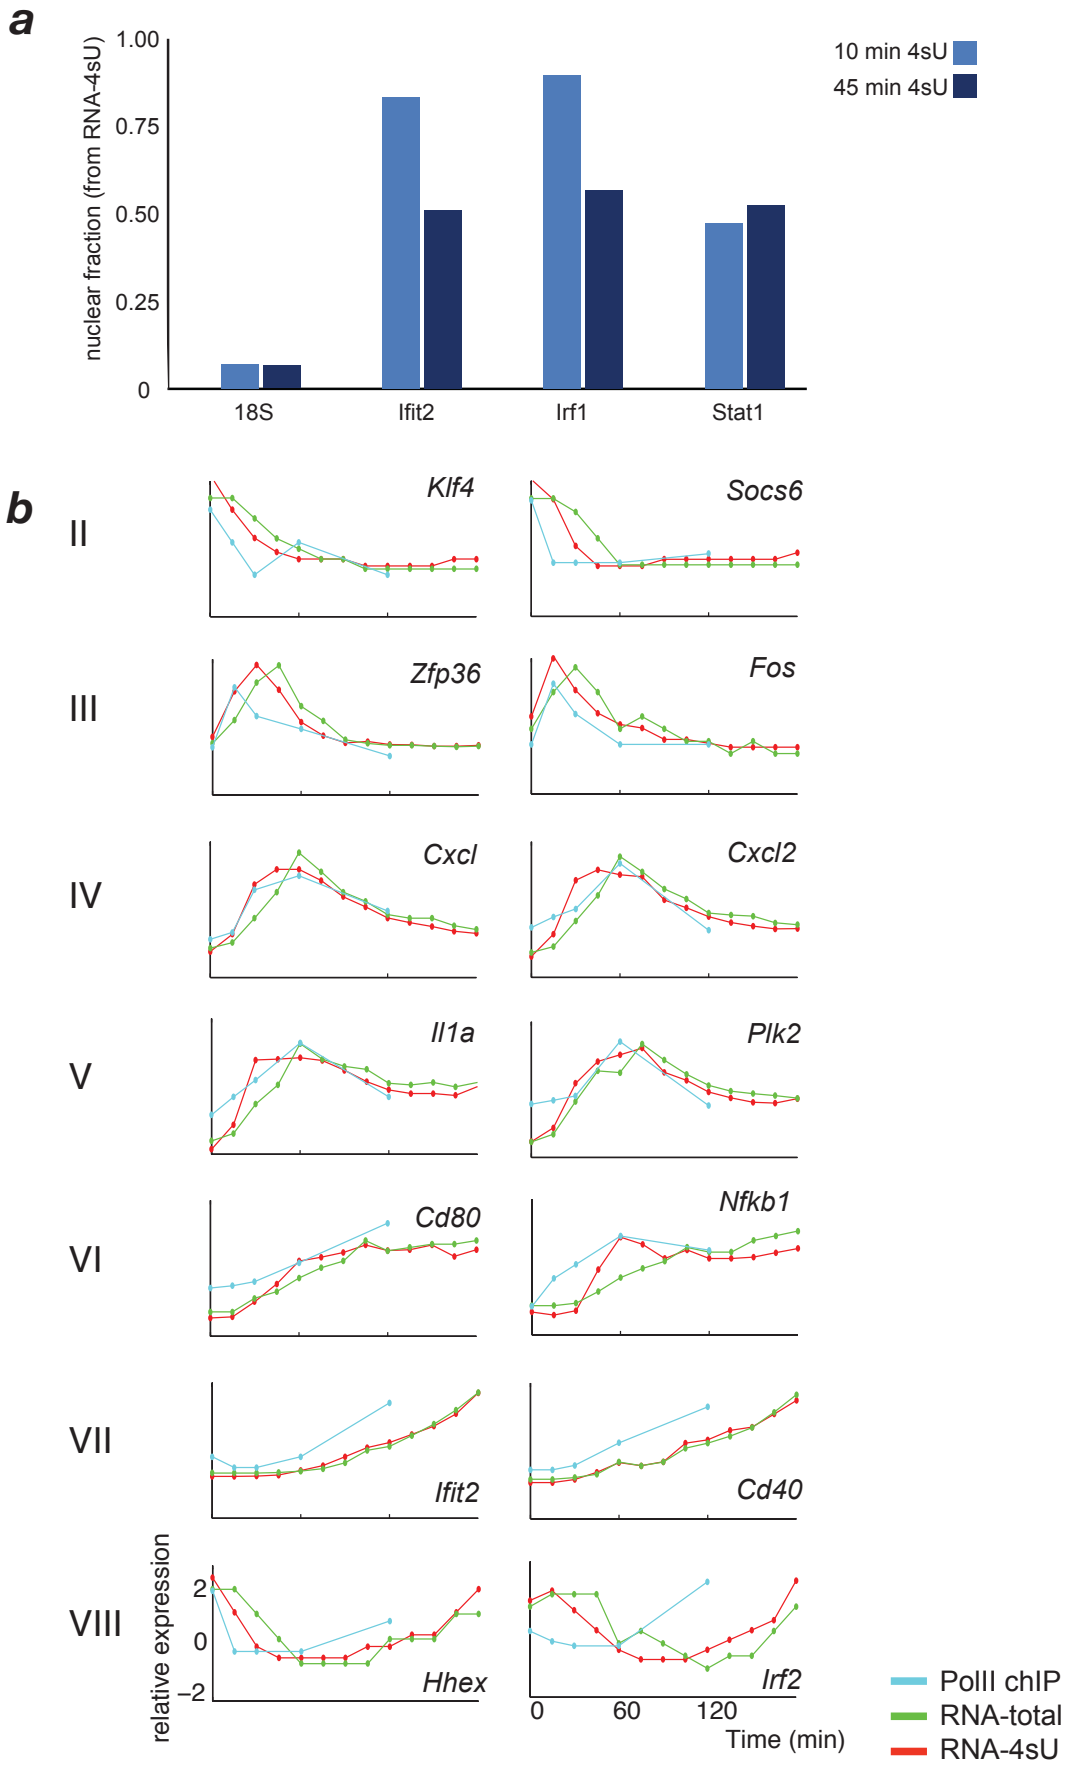
a
nuclear fraction (from RNA-4sU)
1.00
0.75
0.50
0.25
0
18S
Ifit2
Irf1
Stat1
10 min 4sU
45 min 4sU
b
II
III
IV
V
VI
VII
VIII
Klf4
Socs6
Zfp36
Fos
Cxcl
Cxcl2
Il1a
Plk2
Cd80
Nfkb1
Ifit2
Cd40
Hhex
Irf2
relative expression
2
0
−2
0
60
120
Time (min)
PolII chIP
RNA-total
RNA-4sU

Supplementary Figure 6: 4sU-labeled RNA following 10 min of labeling measures transcription rates in unperturbed cells.

(a) After 10 minutes labeling, most 4sU purified RNA is nuclear. Shown is the fraction of nuclear 4sU-RNA (out of nuclear and cytoplasmic 4sU-RNA) from DCs collected at 3h post-LPS stimulation, after different metabolic labeling times (10 and 45 minutes, light and dark blue, respectively). Expression levels in each fraction were quantified using qRT-PCR for 3 induced genes (ifit2, irf1 and stat1) and two controls (28S, 18S). The 28S measurement is used for normalization, and thus not shown. The low nuclear levels of 18S might be due to contamination with unlabeled rRNA, which is more significant for rRNA because of its high abundance in the cell (~98% of cellular RNA). **(b)** RNA-polymerase II occupancy at promoters of LPS-induced genes agrees with RNA-4sU measurements following 10 min labeling. Shown are relative pol-II ChIP enrichment values (Y axis, blue curve), normalized relative to control genes (*Crytalin, b-globin*), at 6 time points post LPS stimulation (X axis, 0h, 15, 30 min, 1 and 3 hours) for two representative genes from each of the clusters II-VIII from **Fig. 2** (denoted on left). The relative RNA-4sU (red) and RNA-total (green) levels of each gene are also shown for reference.

Supplementary Figure 7

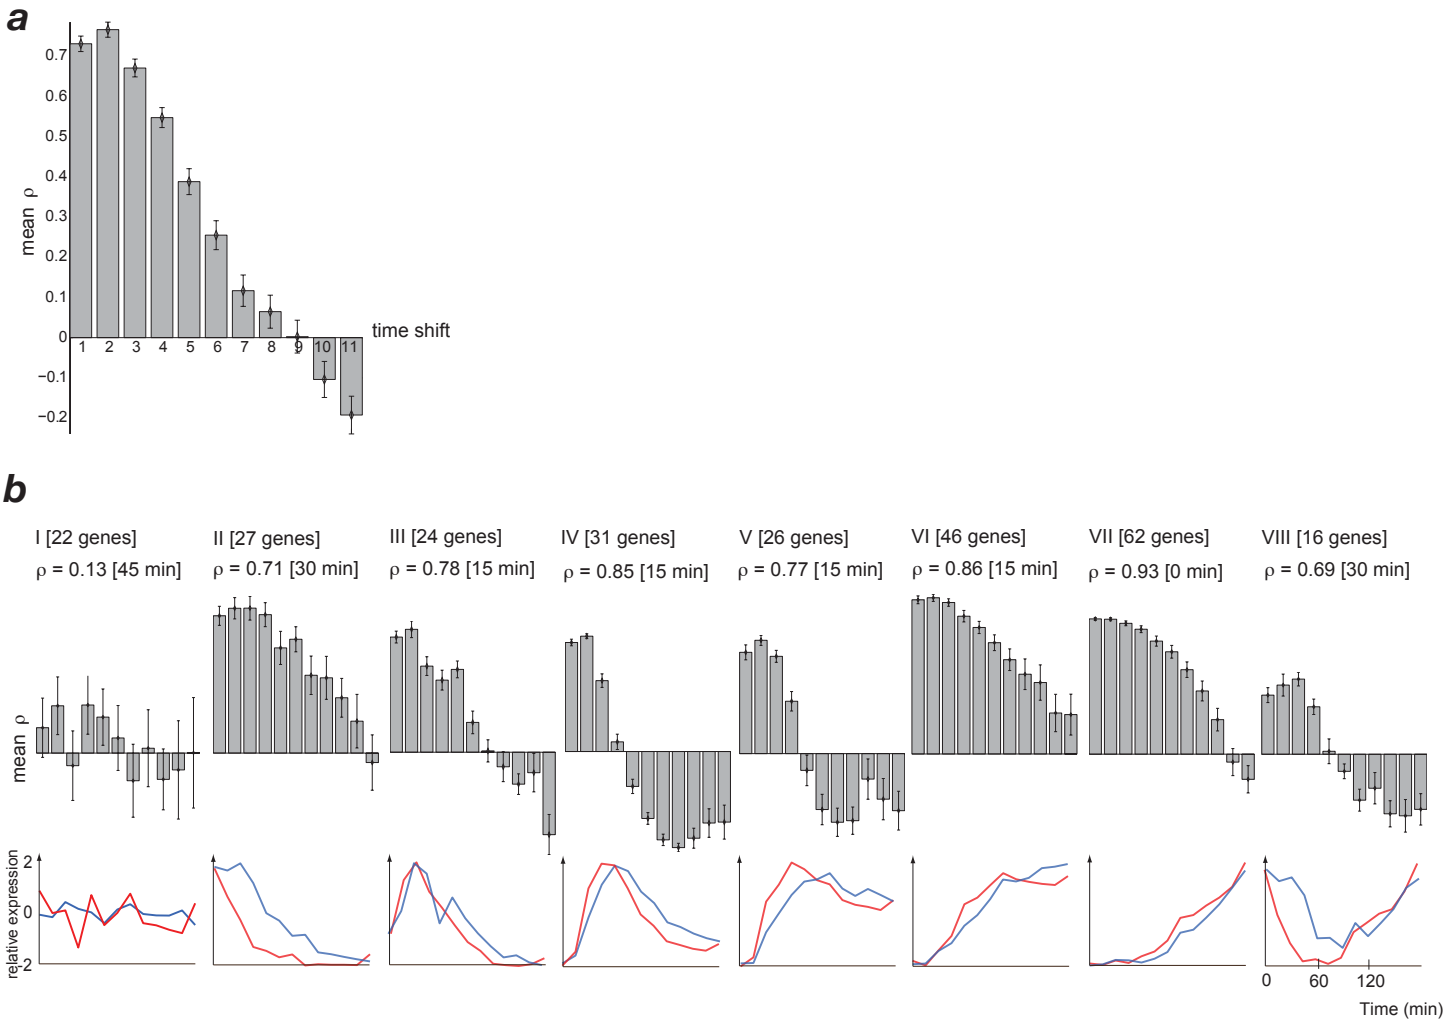

.

Supplementary Figure 7: Peaks in transcription precede peaks in expression by 15-30 minutes.

(a) Shown is the average time-shifted correlation coefficient (ρ) between RNA-4sU and RNA-total profiles for each of the 11 possible time shifts (1=no shift, 2= one time point (15 min) shift, etc.) across all 254 signature genes. **(b)** Top: same as (a), but for genes in each cluster I-VIII separately (clusters numbered as in **Fig. 2,** brackets: number of genes in cluster). Pearson correlation coefficient (ρ) of the best time-lag correlation between the two profiles is indicated on top, with the optimal time lag in brackets. Bottom: the cluster-averaged profile for RNA-4sU (red) and RNA-Total (blue) at each time point.

Supplementary Figure 8

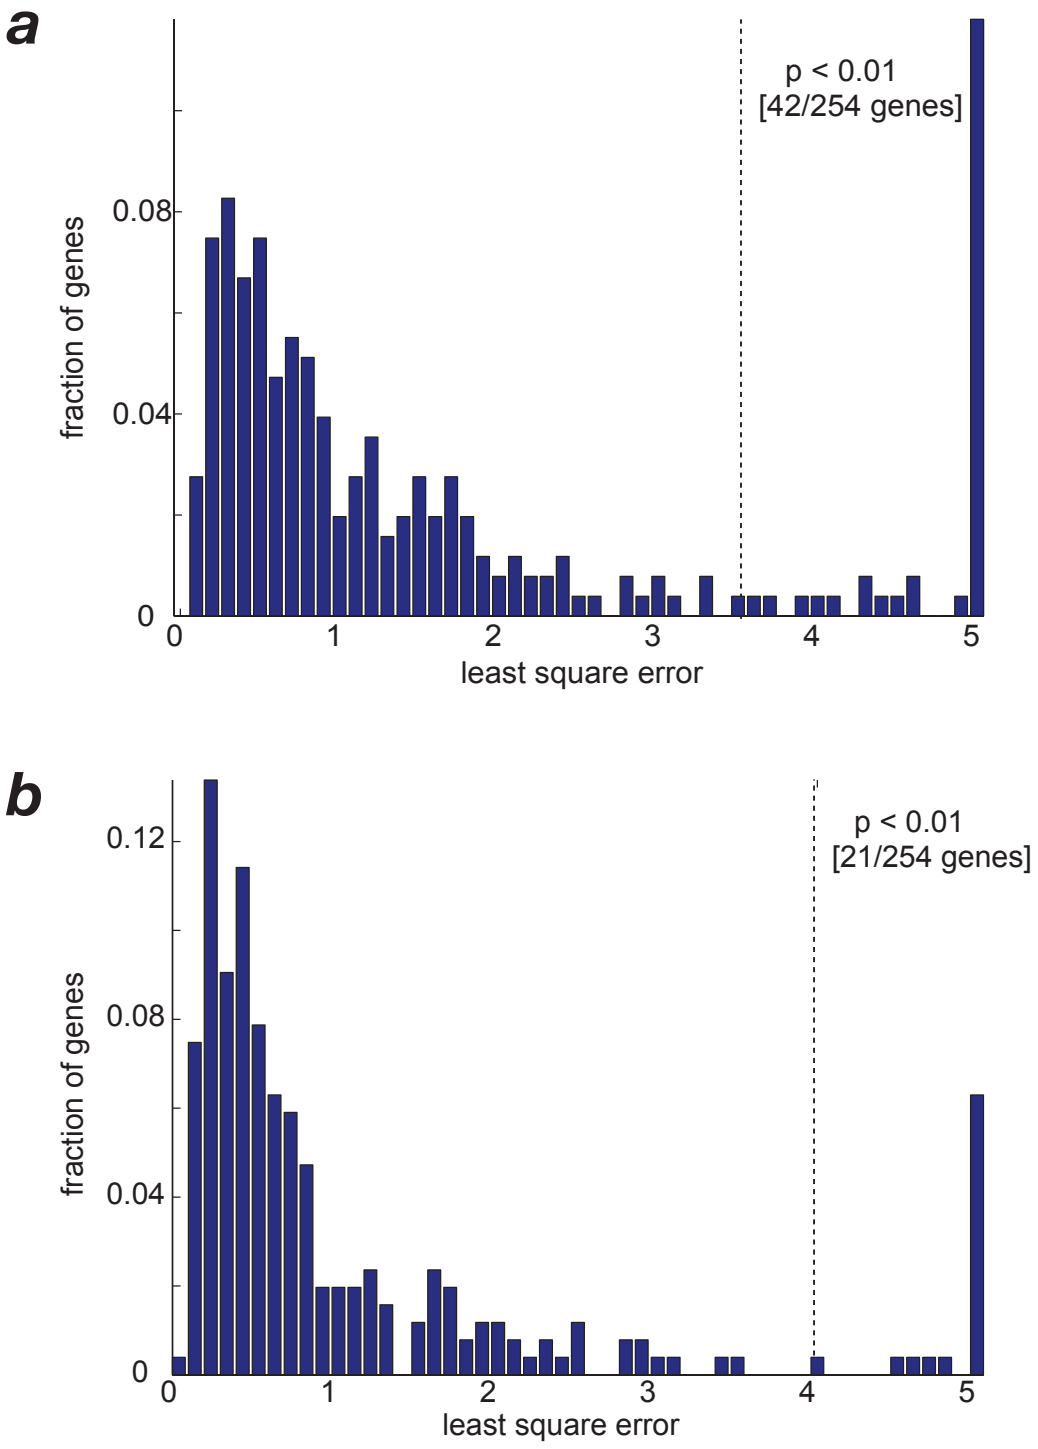

.

Supplementary Figure 8: 'Goodness-of-fit' tests for the constant and varying degradation models

(a) The 'constant degradation' model fits the majority of genes well. Shown is the distribution of the least-square error of the 'constant degradation' model. Dashed line: the threshold for rejecting constant degradation ($p<0.01$): 16% (42/254) of the cases are rejected. **(b)** The 'varying degradation' model improves the fit of the data. Shown is the distribution of the least-square error of the 'varying degradation' model across all 254 genes in the 'signature set'. Dashed line: the threshold for rejecting varying degradation ($p<0.01$): only 8% (21/254) of the cases are rejected.

Supplementary Figure 9

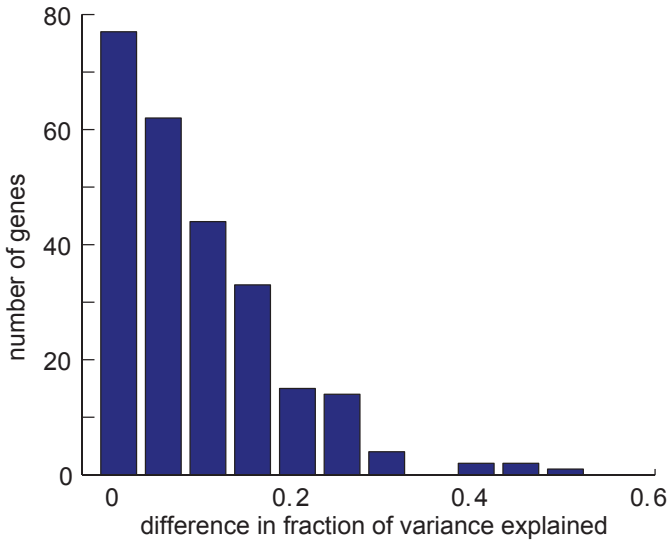

.

Supplementary Figure 9: The ‘varying degradation’ model does not dramatically improve the fraction of explained variation in the data compared to the ‘constant degradation’ model.

Shown is the distribution of the difference in the fraction of variance explained for each gene by the ‘constant degradation’ and the ‘varying degradation’ model. Positive values are the extra fraction explained by varying degradation.

Supplementary Figure 10

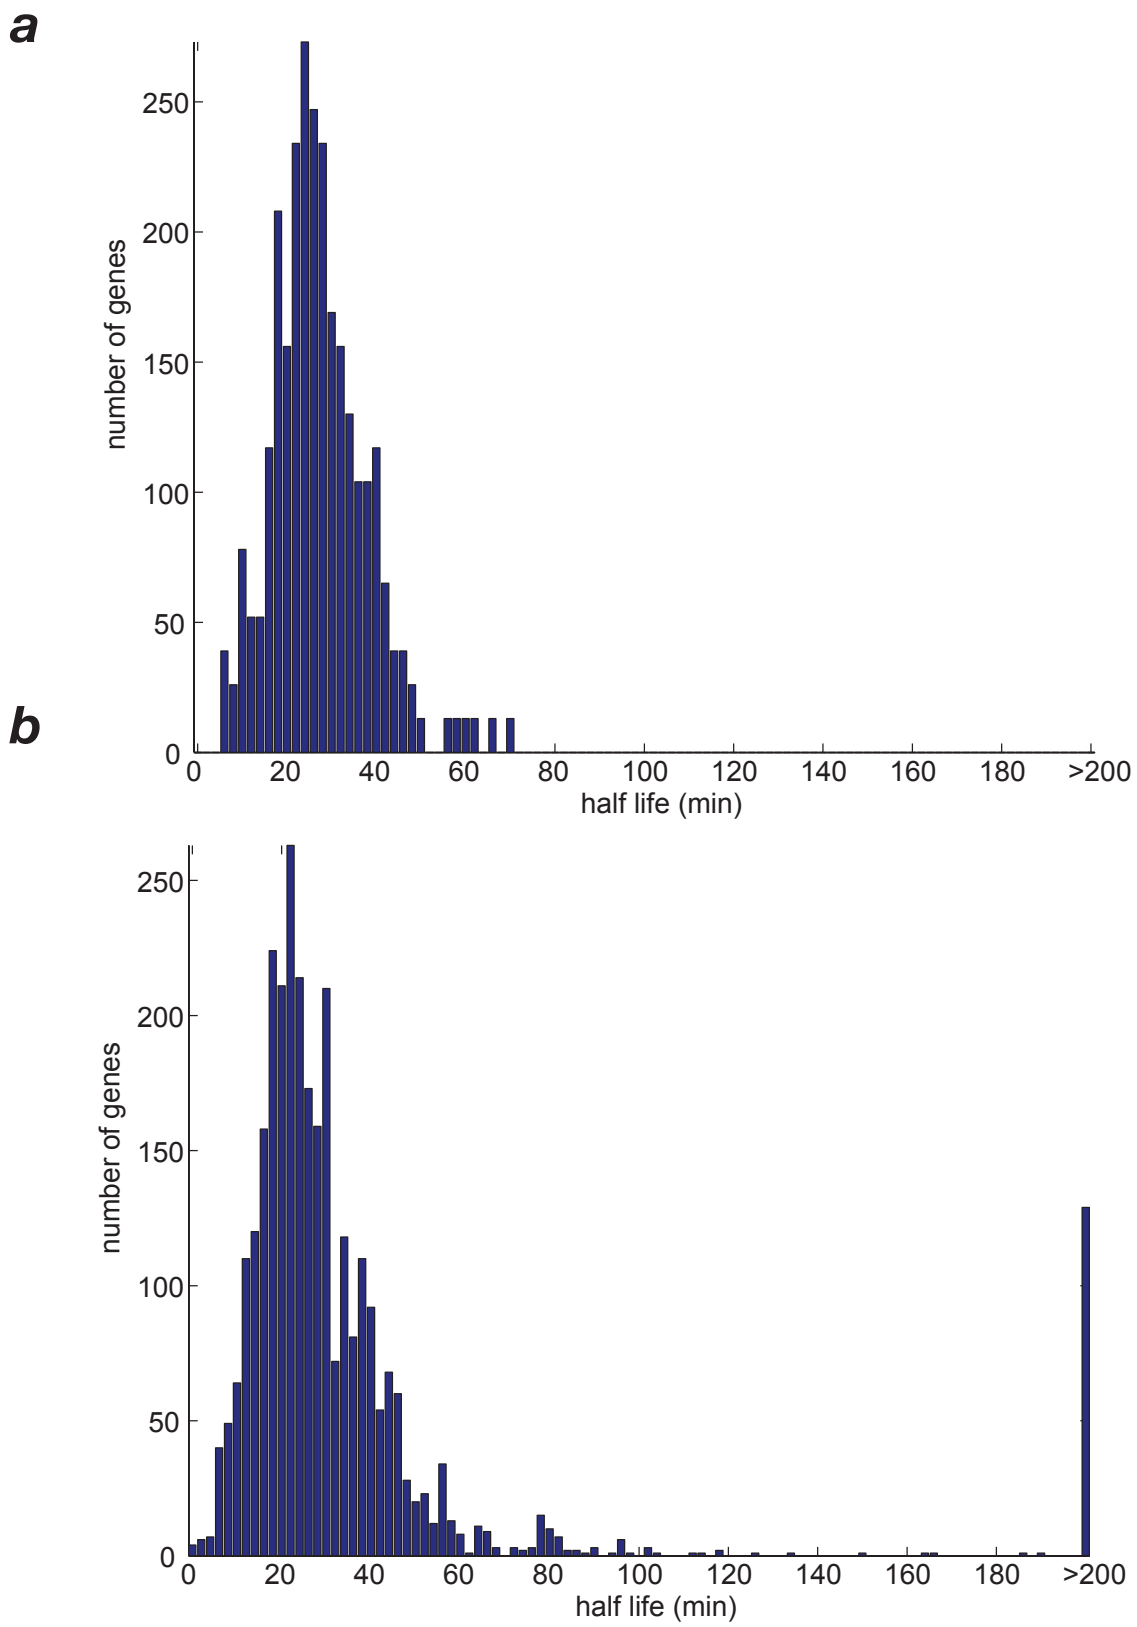

Supplementary Figure 10: Distribution of RNA half-lives predicted by the 'constant' and 'varying' degradation models.

(a) The distribution of RNA half-lives (minutes) estimated by the 'constant degradation' model. **(b)** The distribution of RNA half-lives (minutes) estimated by the 'varying degradation' model at each of the 13 time points in which RNA time-course measurements were taken (between 0 to 3 hours, in 15 minutes intervals).

Supplementary Figure 11

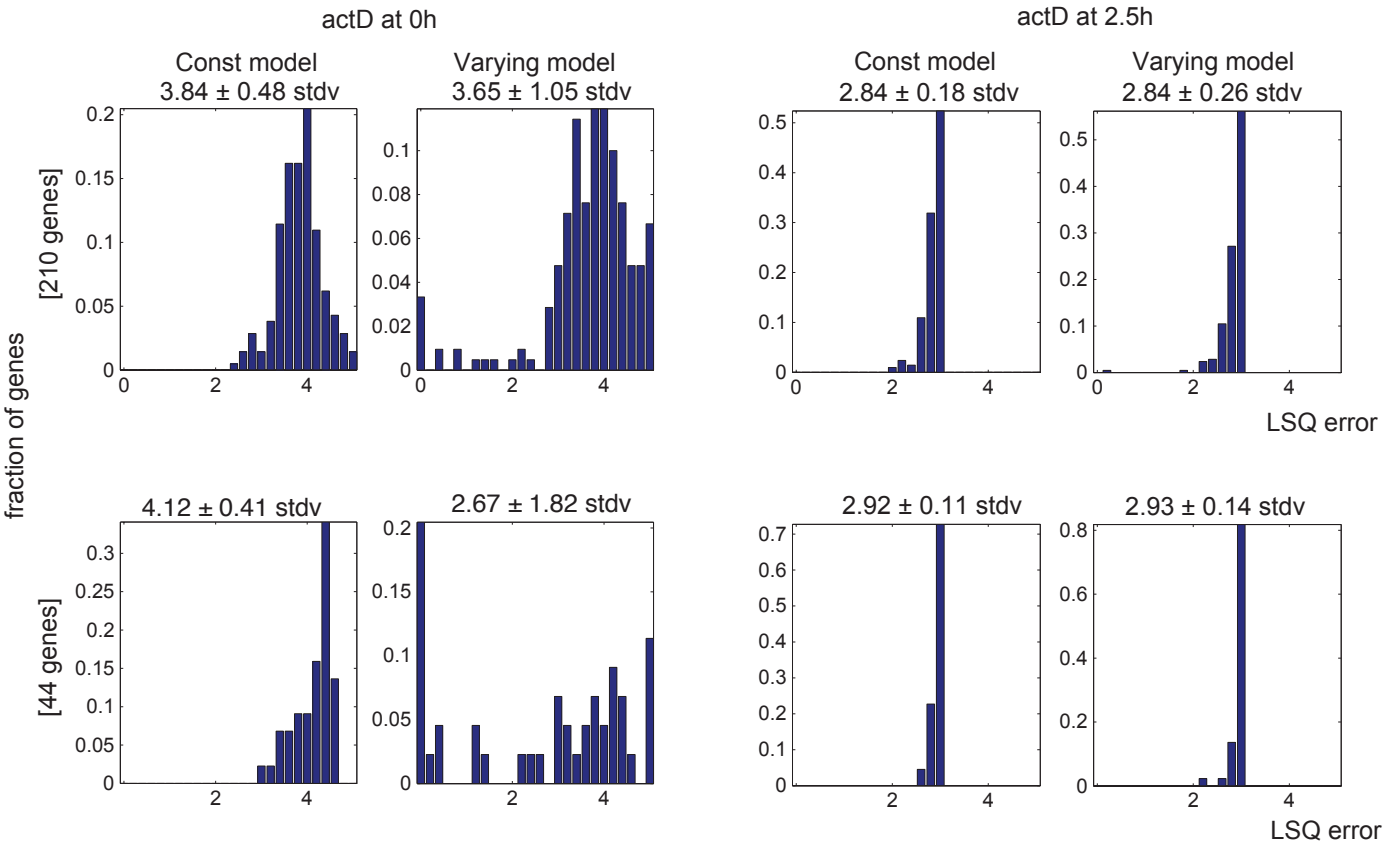

.

Supplementary Figure 11: Comparing the Actinomycin D data with the model predictions.

Comparison of the fit of predictions by the 'constant degradation' and 'varying degradation' models to RNA levels measured following Actinomycin D treatment. Shown are distributions of least-square-error (LSQ, X-axis; normalized by expression) between RNA levels measured at several time points after the addition of Actinomycin D and predicted levels based on either the constant (left) or varying (right) models, for the 210 genes that accept the 'constant degradation' hypothesis (upper row) and the 44 genes that reject it (bottom row). Actinomycin D was added either before LPS treatment (actD at 0h; two left panels) or at 2.5 hours after LPS treatment (actD at 2.5h; two right panels). The model predictions of RNA levels are based on the RNA levels at the time of Actinomycin D addition, on the model (either 'varying' or 'constant') degradation rate at the time of Actinomycin D addition and assuming transcription arrest after Actinomycin D addition. The mean and standard-deviation of LSQ scores is indicated at the top of each histogram. Compared to RNA levels measured following Actinomycin D treatment before LPS stimulation (0h, two left panels), the 'varying degradation' model improves the fit of the 44 genes that reject the 'constant degradation' model (bottom), and has comparable fit as the 'constant degradation' model for the remaining genes (top). The two models have comparably good fit for all genes to RNA levels measured following Actinomycin D treatment at 2.5h post-LPS (two right panels).

Supplementary Figure 12

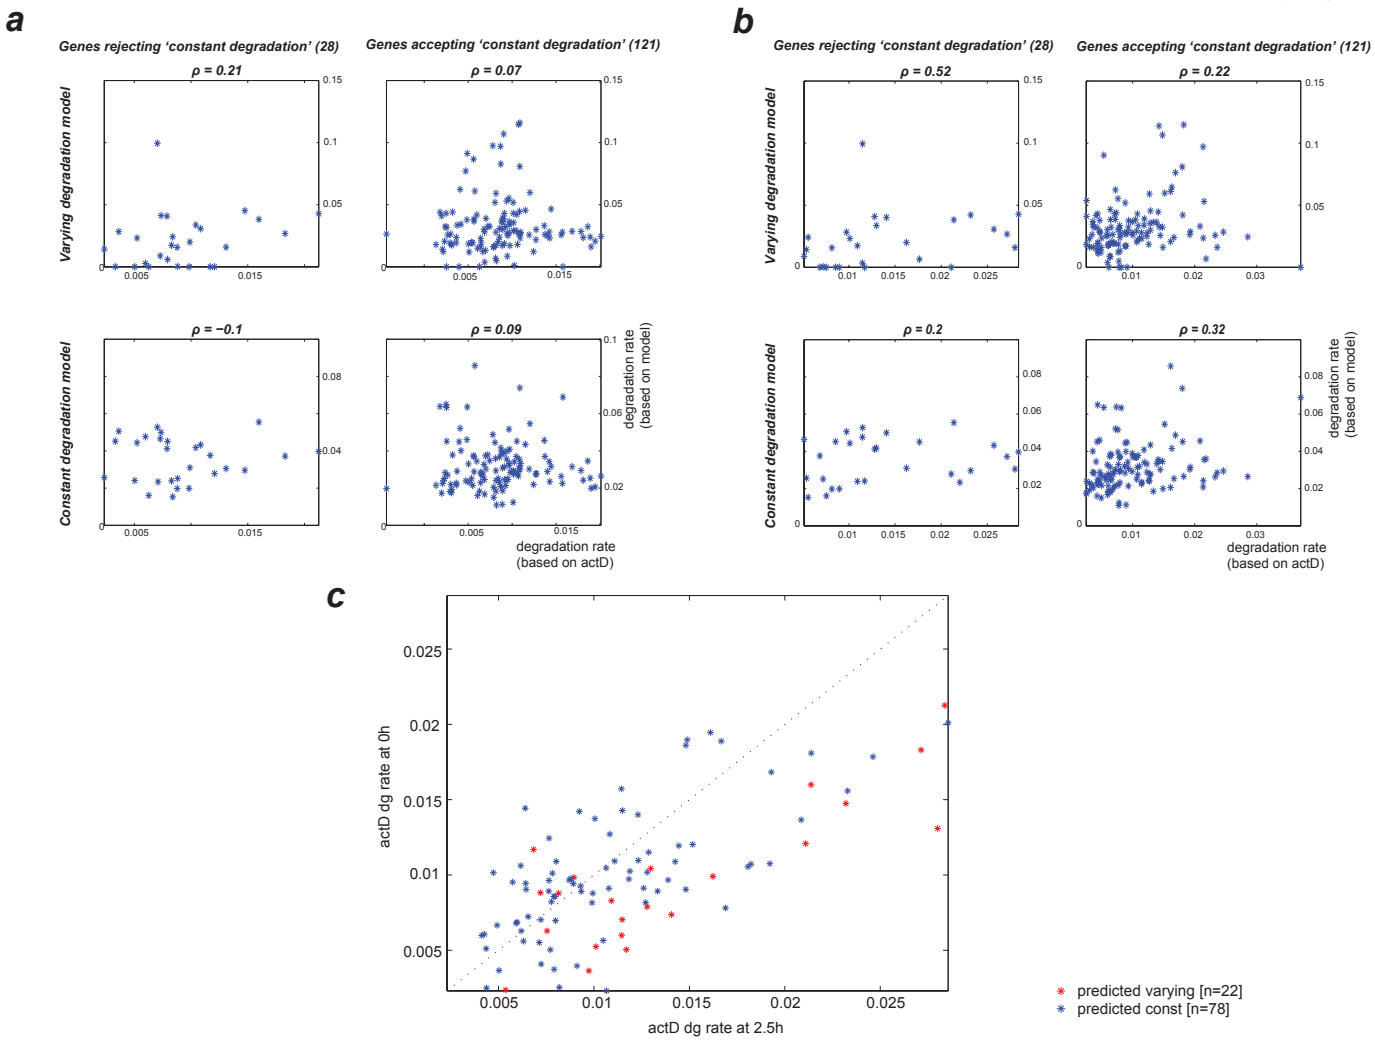

Supplementary Figure 12: Comparing RNA degradation rates estimated by the Actinomycin D data and the 'varying' and 'constant' degradation models.

(a-b) The genes that reject the 'constant degradation' hypothesis correlate better with the 'varying degradation' model predictions, while for the accepting genes the correlation with the 'constant degradation' model is higher. Shown is the correlation coefficient (spearman ρ; top of each plot) between the estimated degradation rates using standard Actinomycin D protocols (X axis; **Methods**) and the predicted degradation rate based on the models (Y axis; varying degradation model – top row, constant degradation model – lower row), for genes that reject the 'constant degradation' model predictions (left column) compared with genes that do not (right column). Only genes for which Actinomycin D rate is estimated with r-squared of 0.6 or higher are included in the analysis (149/254). Shown in **(a)** are the results when adding Actinomycin D at 0 hours, and in **(b)** the results when adding Actinomycin D at 2.5 hours after LPS stimulation. **(c)** Genes that reject the 'constant degradation' model (red) show a bigger shift between rates predicted based on Actinomycin D data before LPS treatment (Y-axis) and at 2.5 hours after LPS stimulation (X-axis) compared with genes that retain the 'constant degradation' model (blue). Most commonly, degradation rates for these genes increase at 2.5 hours. Only genes with r^2 of 0.8 or higher for predicted degradation rate are considered (100/254 genes).

Supplementary Figure 13

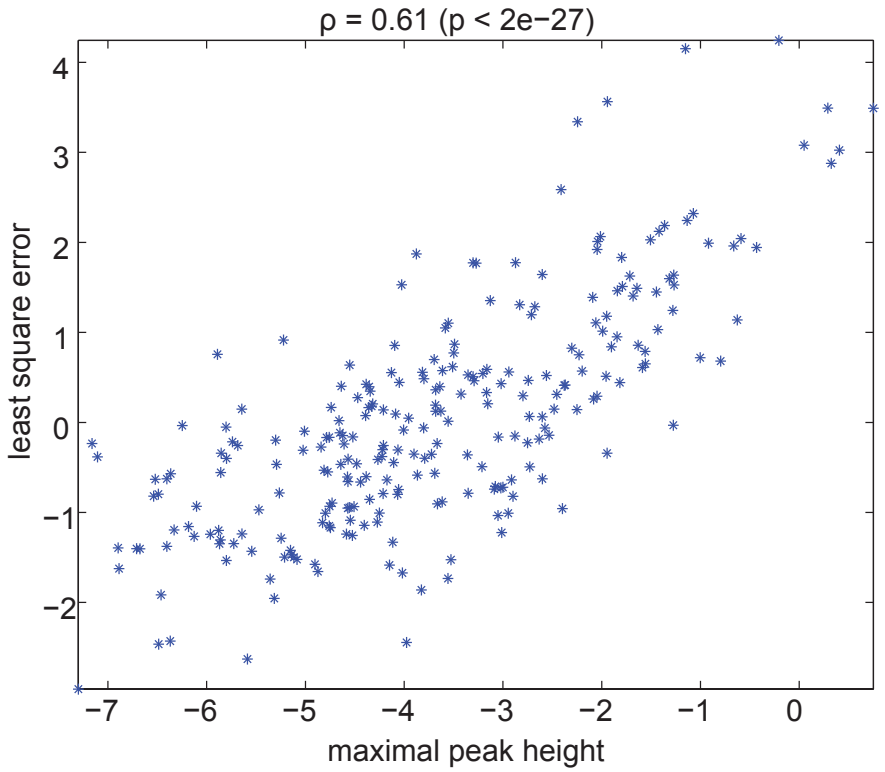

.

Supplementary Figure 13: The ‘constant degradation’ model fails most significantly at peaks in expression.

Shown is the correspondence between the least squared error in the ‘constant degradation’ model (Y axis, log scale) and the height of the maximal peak in the expression profile (X axis, log scale). The Pearson correlation coefficient (ρ) is indicated on top, along with the statistical significance of the correlation.

Supplementary Figure 14

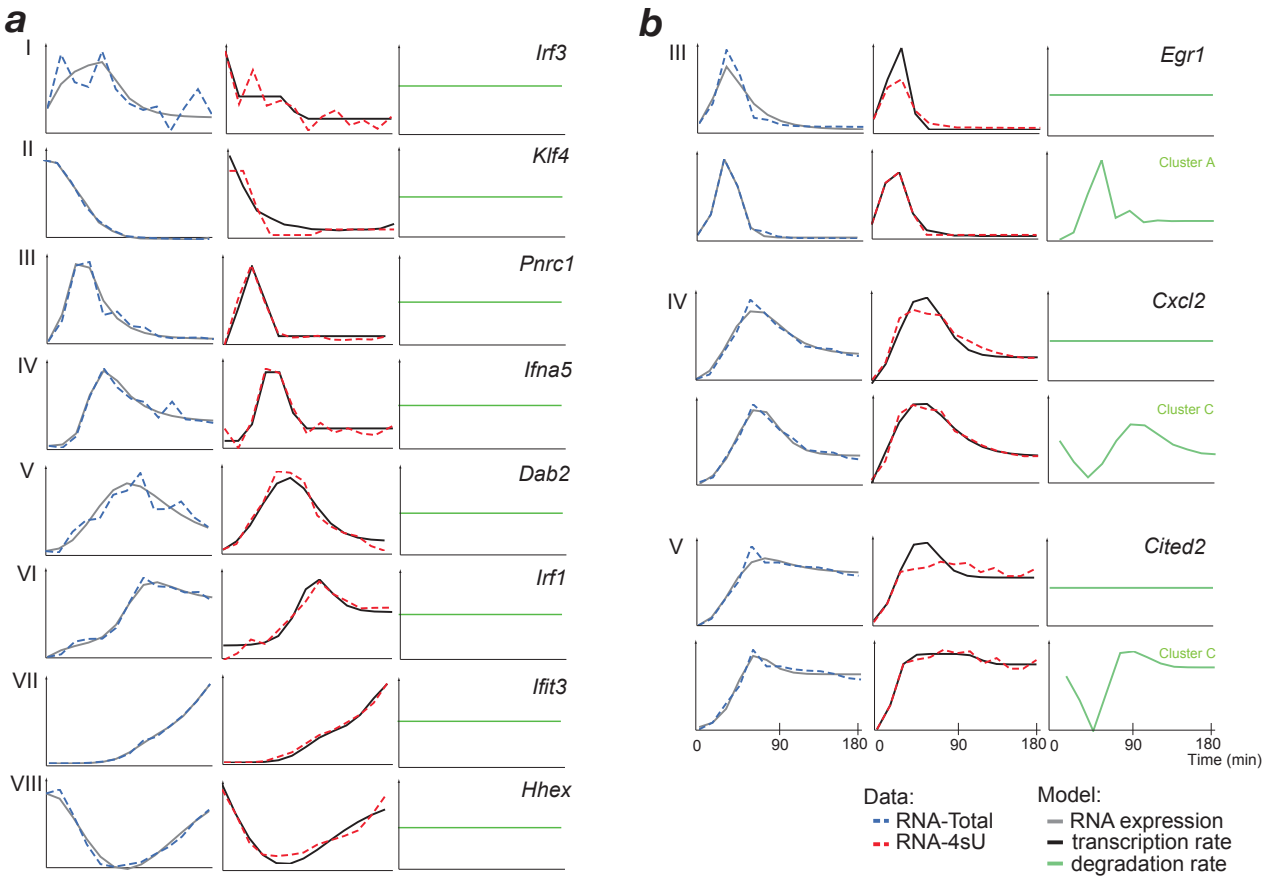

Supplementary Figure 14: Constant degradation rates explain most of the expression dynamics observed, while changes in degradation rates help shape peaks. **(a)** The 'constant degradation' model explains most of the expression dynamic observed in the data. Shown are illustrative examples of eight genes (one per cluster I-VIII; Left: name and cluster assignment) that accept the 'constant degradation' model. Gray, black and green lines: model predictions (expression, transcription rate, and degradation rate, respectively); dashed blue and dashed red lines: data (RNA-total and RNA-4sU, respectively); Left: expression level; middle: transcription rate; right: degradation rate (estimate only). **(b)** Genes with peaked responses reject the 'constant degradation' model. Shown are examples of three genes and their fit to the 'constant degradation' model (Upper row) and to the 'varying degradation' model (lower row). Panels formatted as in (a). For 'varying degradation' predictions, the degradation cluster (A, B or C) is indicated at the top of the degradation profile.

Supplementary Figure 15

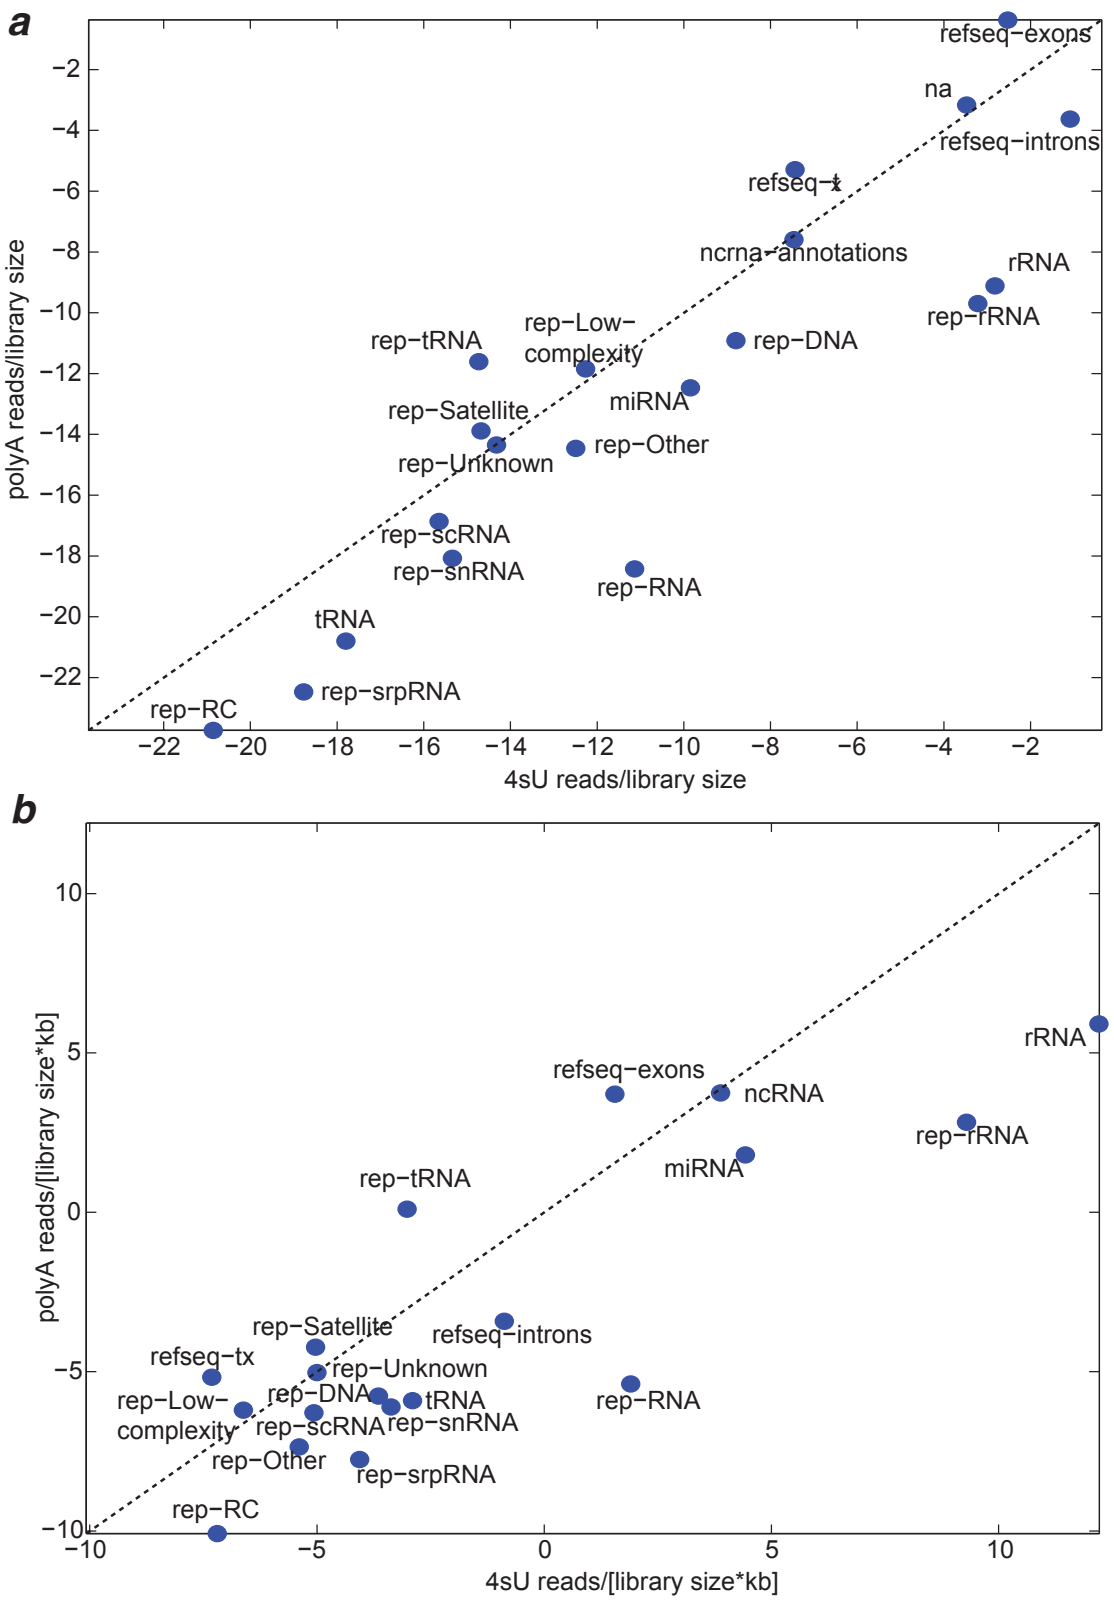

Supplementary Figure 15: 4sU-Seq captures a broader range of transcripts compared to polyA+ RNA-Seq.

Shown is the fraction of reads in RNA-4sU-Seq libraries (x-axis, log scale) and polyA+ RNA-Seq libraries (y-axis, log scale), for several annotation categories. Only reads that were mapped to a unique location in the genome or to rRNA are considered. **(a)** Reads in each of the libraries are normalized for library size only. **(b)** Reads in each of the libraries are normalized for both library size, and the overall size of the annotation category (in Kb).

Supplementary Figure 16

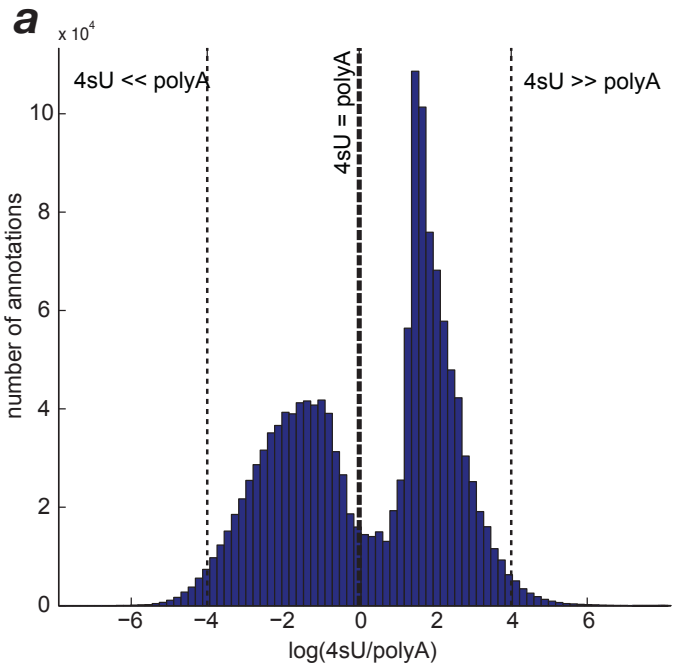

b

	4sU << polyA	4sU >> polyA
mir_annotations	0.01	0.03
not annotated	14.49	10.58
ncrna_annotations	0.67	0.70
refseq_exons	77.37	0.31
refseq_introns	5.09	87.09
refseq_tx	2.34	0.00
rep_DNA	0.00	0.06
rep_Low_complexity	0.02	0.01
rep_Other	0.00	0.07
rep_RNA	0.00	0.03
rep_Satellite	0.00	0.00
rep_Unknown	0.00	0.00
rep_rRNA	0.00	0.74
rep_scRNA	0.00	0.00
rep_snRNA	0.00	0.03
rep_srpRNA	0.00	0.00
rep_tRNA	0.00	0.01
rrna_annotations	0.00	0.33
trna_annotations	0.00	0.01

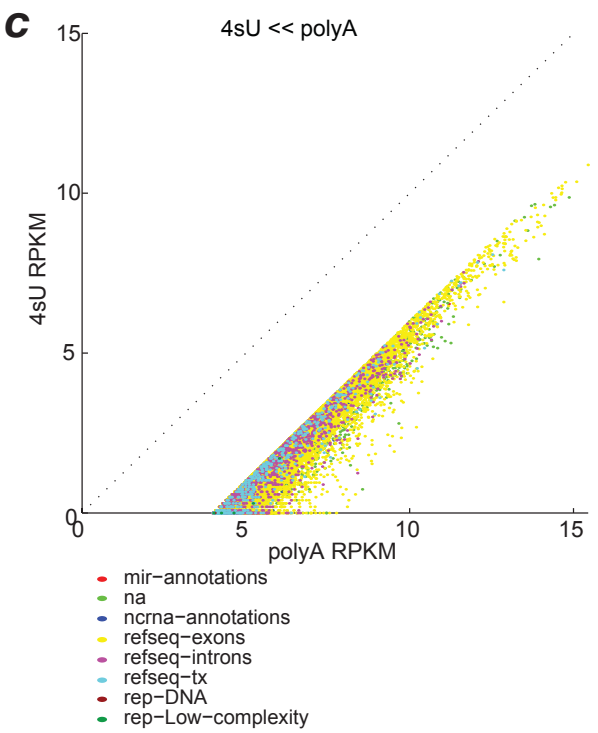

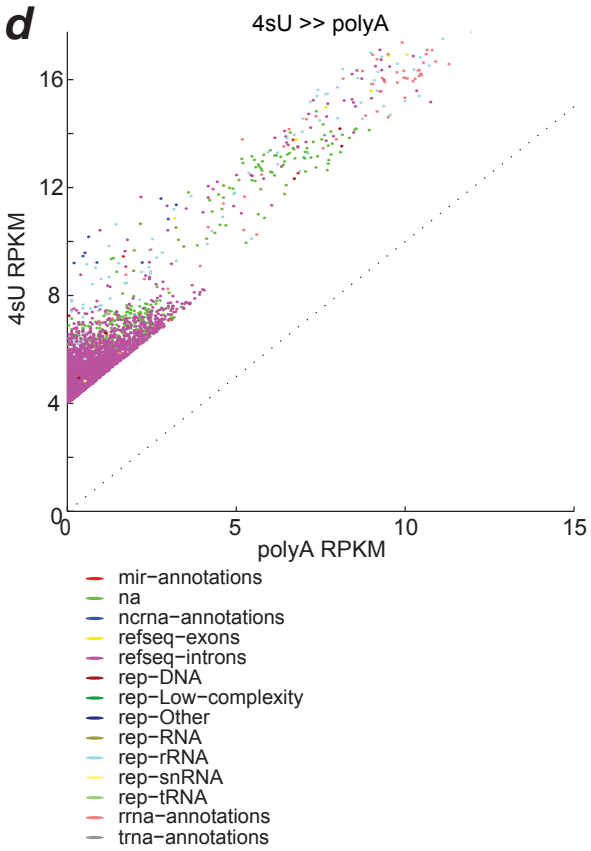

Supplementary Figure 16: Introns are enriched in 4sU-Seq reads, while exons are enriched in RNA-Seq reads.

We tiled the entire mouse genome with 50bp tiles (not overlapping) and counted the number of reads that mapped to each fragment in RNA-Seq vs. 4sU-Seq libraries. **(a)** Distribution of the ratio of reads in 4sU-Seq vs. RNA-Seq per genomic tile. Bold dashed line: 4sU-Seq reads are equal to RNA-Seq reads. Right and left light dashed line: tiles with at least 4 times more reads in 4sU-Seq (4sU-Seq enriched) or in RNA-Seq, (RNA-Seq enriched) respectively. **(b)** RNA-Seq enriched tiles are exonic, 4sU-Seq enriched tiles are intronic. The percentage of RNA-Seq enriched tiles (middle column) or 4sU-Seq enriched tiles (right column) in several annotation categories (rows). Bold (red): the most common annotation: 77% of RNA-Seq enriched tiles are exonic; 87% of 4sU-Seq enriched tiles are intronic. **(c,d)** Distributions of reads in RNA-Seq enriched tiles (c) or 4sU-Seq enriched tiles (**d**) in 4sU-Seq libraries (x-axis) vs. RNA-Seq libraries (y-axis). Differently annotated tiles are color-coded as indicated on the graph.

Supplementary Figure 17

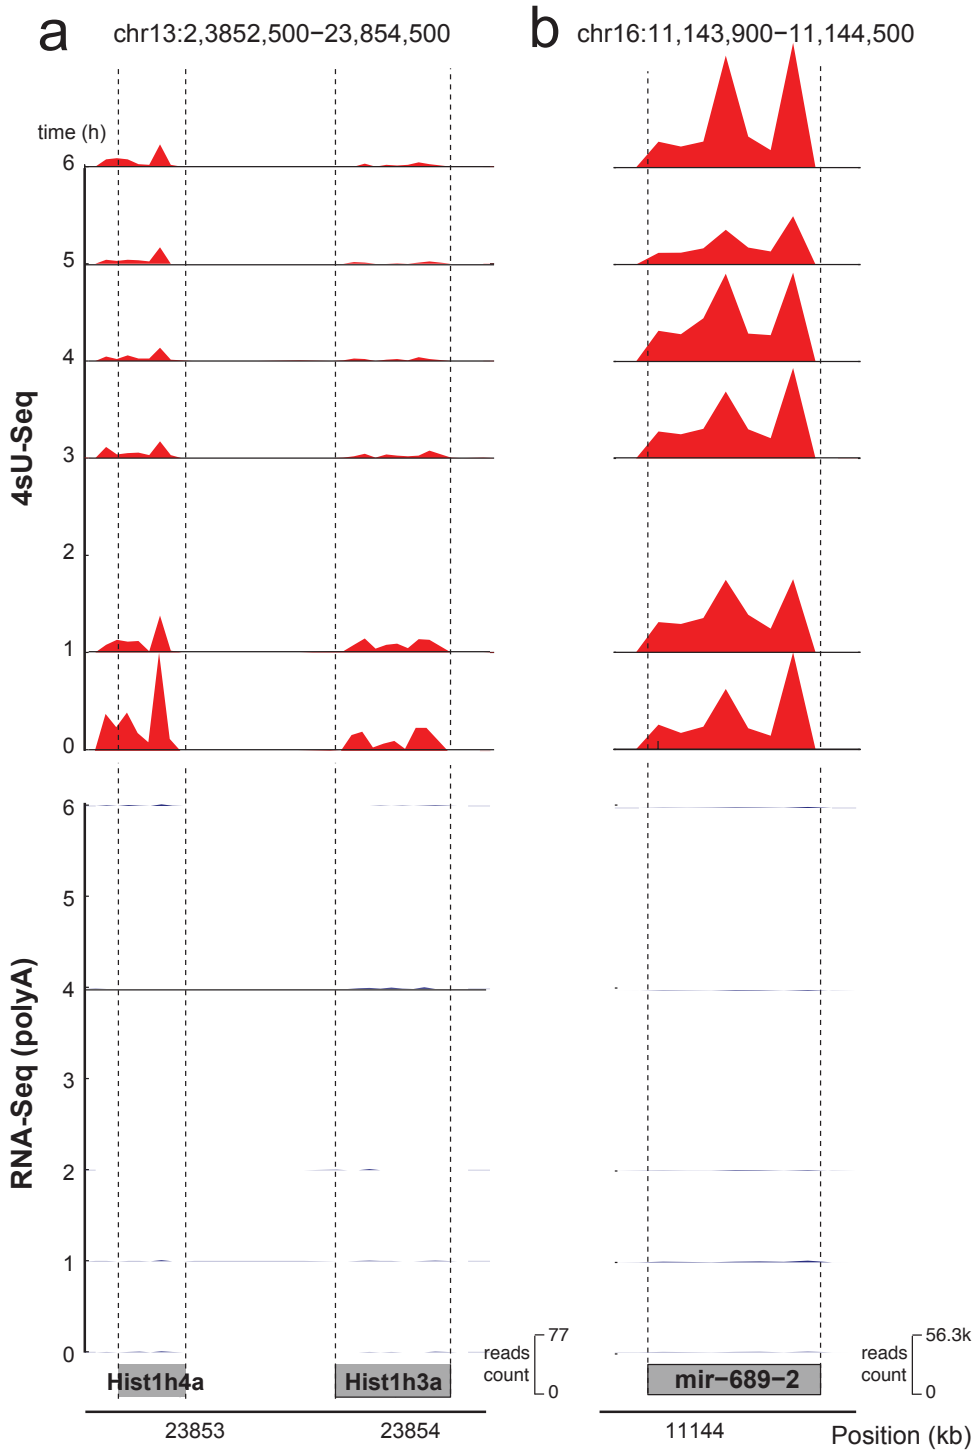

Supplementary Figure 17: Histones and pre-miRNA transcripts are captured by 4sU-Seq but not by RNA-Seq.

Shown is the read coverage (height) at each gene at each time point (row) in each library (top: 4sU-Seq, red; bottom: RNA-Seq, blue) for two neighboring histone genes **(a)** and for a miRNA gene **(b)**.

Supplementary Figure 18

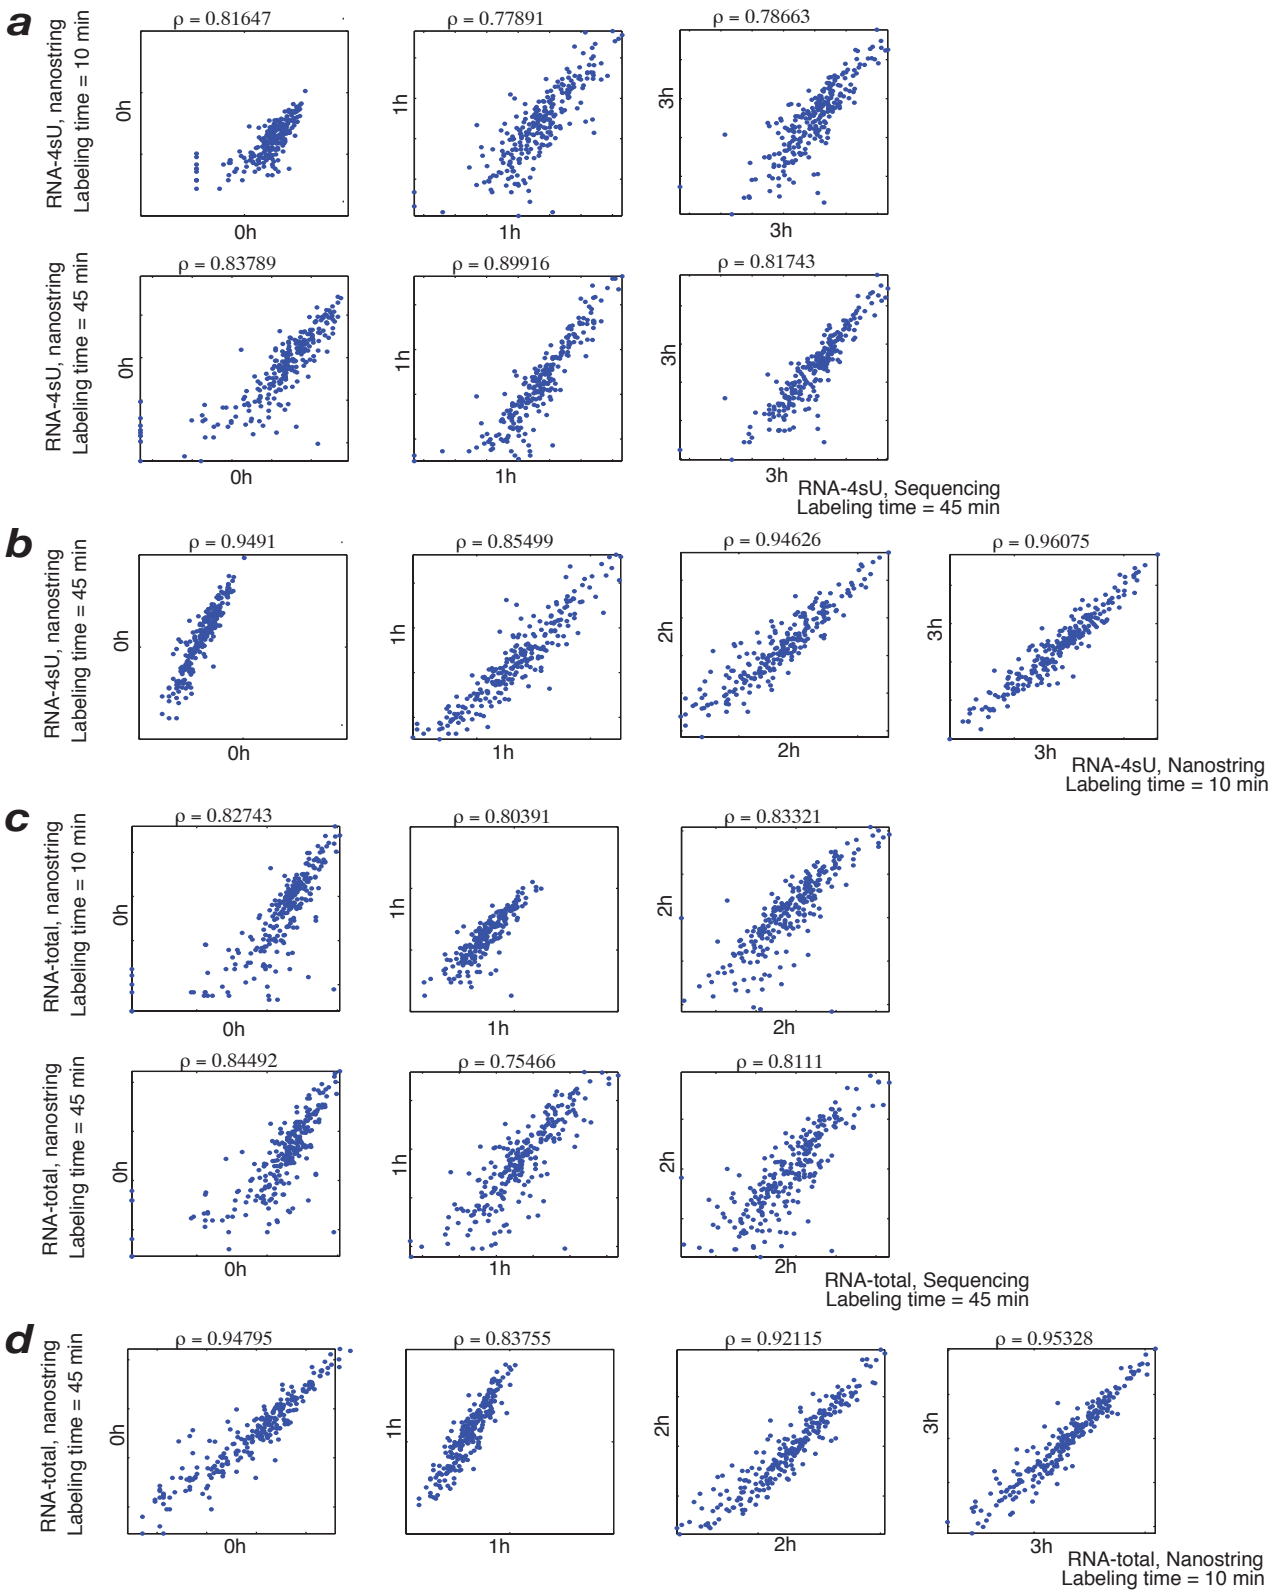

Supplementary Figure 18: RNA level estimates based on RNA-Sequencing and nCounter are highly correlated.

Shown is the correspondence between RNA levels estimated based on nCounter data and RNA sequencing data for the 254 signature set genes. **(a)** Correlation of 4sU labeled RNA measurements, at several times post-LPS stimulation (0, 1, 2 and 3 hours). Sequenced samples (X axis, log scale): 45 min. metabolic labeling times. nCounter samples (Y axis, log scale): 10 min labeling (top) or 45 min labeling (bottom). Pearson correlation coefficient (ρ) is indicated on top. **(b)** As in (a) but comparing nCounter RNA-4sU samples at 10 min labeling (X-axis, log scale) to nCounter RNA-4sU at 45 min labeling (Y-axis, log scale). **(c)** As in (a) but comparing RNA-Total sequenced samples (X-axis, log scale) to nCounter RNA-Total samples (Y-axis, log scale). **(d)** As in (a) but comparing nCounter RNA-Total samples at 10 min labeling (X-axis, log scale) to nCounter RNA-Total at 45 min labeling (Y-axis, log scale).

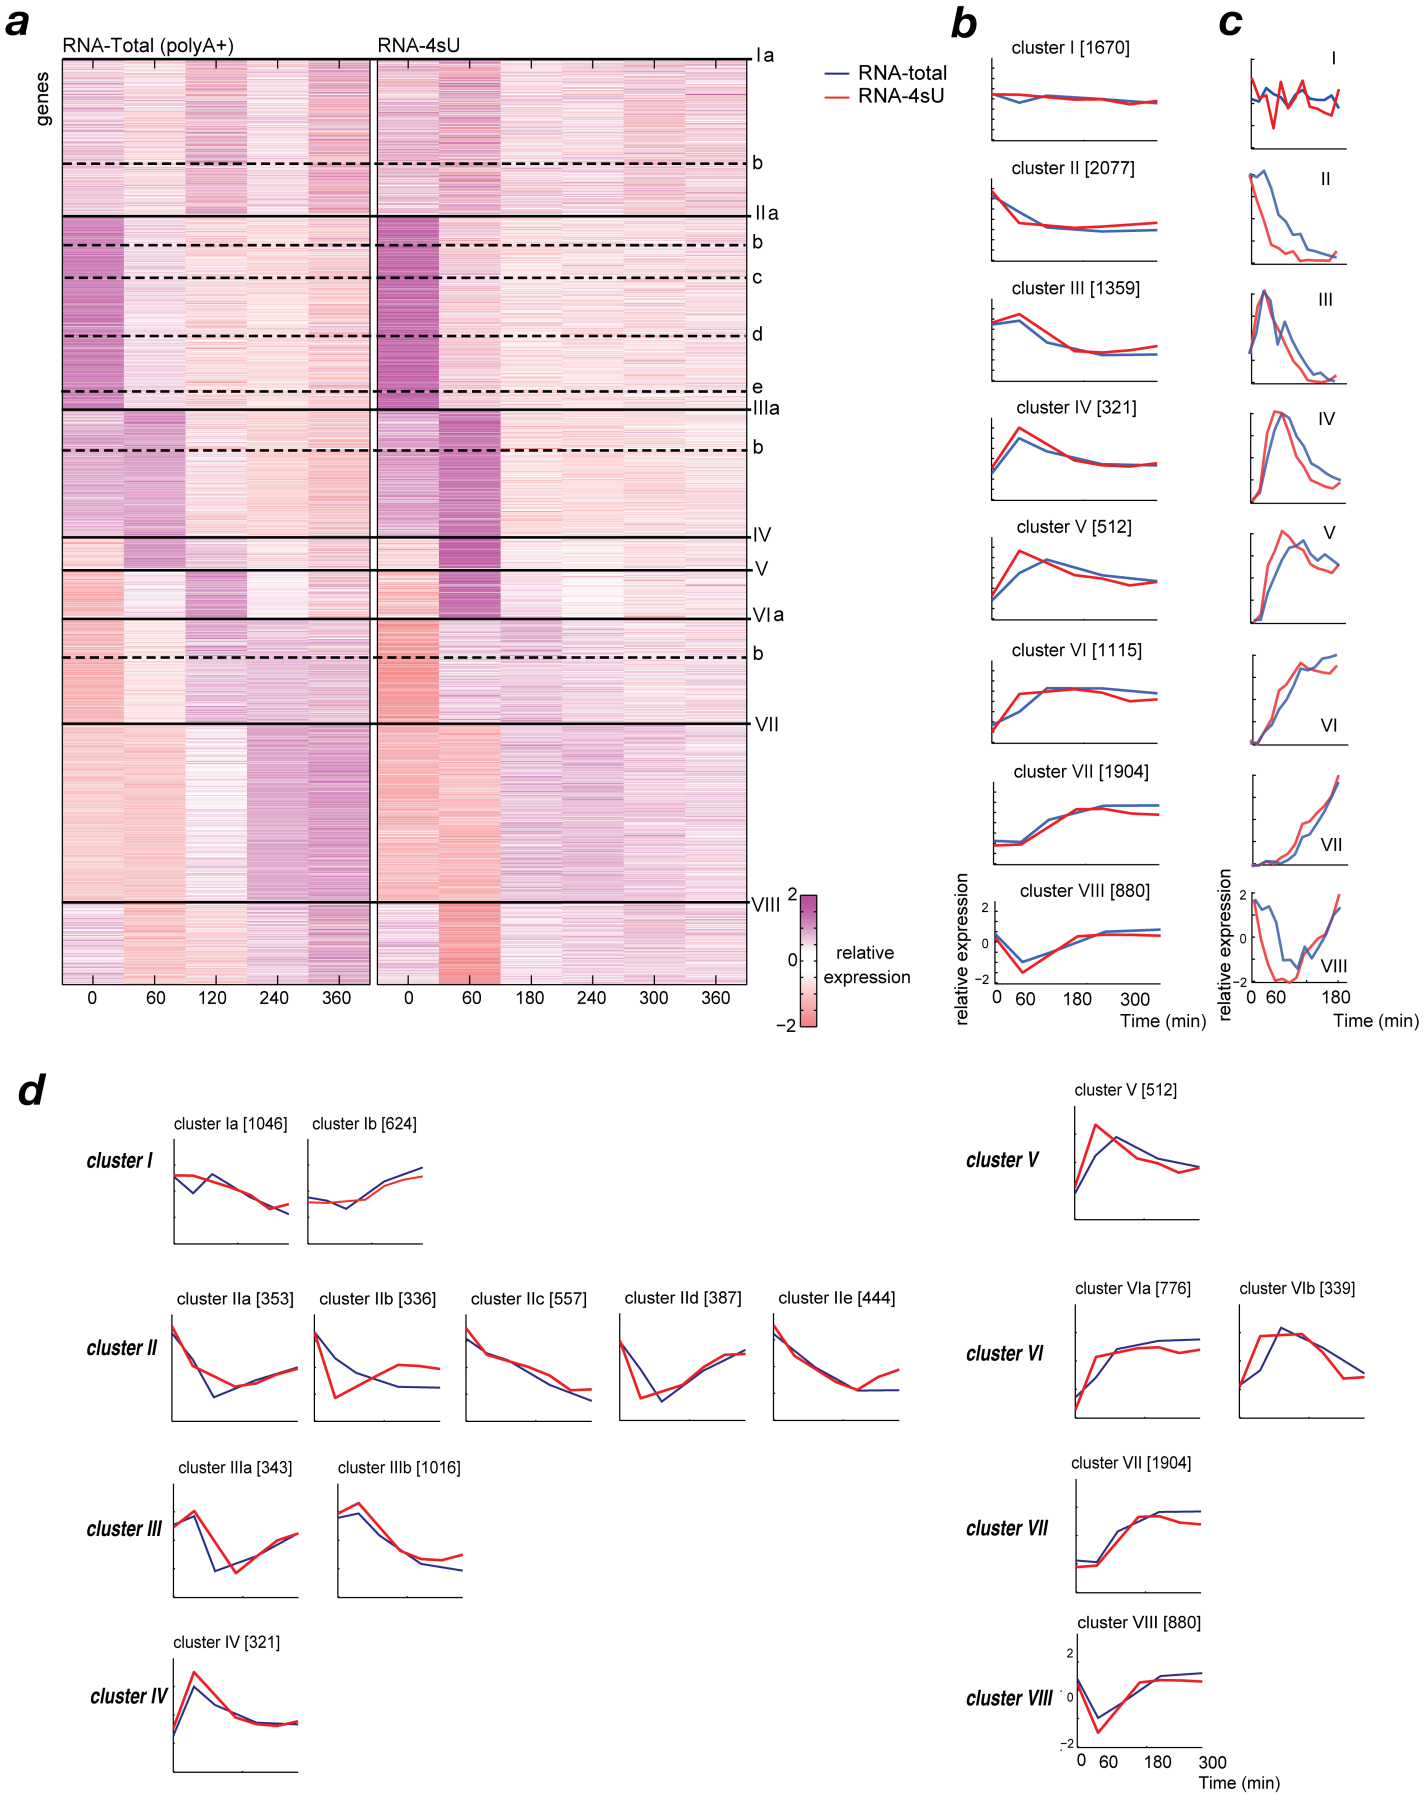
a
RNA-Total (polyA+)
RNA-4sU
genes
Ia
b
IIa
b
c
d
e
IIIa
b
IV
V
VIa
b
VII
VIII
0 60 120 240 360
0 60 180 240 300 360
RNA-total
RNA-4sU
2
0
−2
relative expression
b
cluster I [1670]
cluster II [2077]
cluster III [1359]
cluster IV [321]
cluster V [512]
cluster VI [1115]
cluster VII [1904]
cluster VIII [880]
relative expression
0 60 180 300
Time (min)
c
I
II
III
IV
V
VI
VII
VIII
relative expression
0 60 180
Time (min)
d
cluster I
cluster Ia [1046]
cluster Ib [624]
cluster II
cluster IIa [353]
cluster IIb [336]
cluster IIc [557]
cluster IId [387]
cluster IIe [444]
cluster III
cluster IIIa [343]
cluster IIIb [1016]
cluster IV
cluster IV [321]
cluster V
cluster V [512]
cluster VI
cluster VIa [776]
cluster VIb [339]
cluster VII
cluster VII [1904]
cluster VIII
cluster VIII [880]
2
0
−2
0 60 180 300
Time (min)

Supplementary Figure 19: Temporal profiles of genome-wide data match and refine the 8 patterns from the signature set.

(a) Assignment of the genome-wide sequencing data to the 8 distinct temporal clusters of newly-transcribed and total RNA identified based on the signature set data. Shown are expression profiles (RPKM) for the 10,106 expressed genes (rows) based on RNA-Total (left) and RNA-4sU (right) measurements across 11 time points (columns). Clusters are numbered on right and separated by solid horizontal black lines. Horizontal dashed lines: further sub-clustering of the data. Purple: high relative expression; white: mean expression; pink: low relative expression. **(b,c)** Average profiles for RNA-4sU (red) and RNA-Total (blue) for each cluster based on the sequencing data (b, 0-6h) and the nCounter data (c, 0-3h). The size of each cluster is indicated in brackets. **(d)** Cluster refinement in the extended time-course data shows distinct behaviors at later time points. Shown is the average cluster profile for RNA-4sU (red) and RNA-Total (blue) in each sub-cluster for every one of the 8 main clusters. Original cluster number is indicated to the left; sub cluster size is indicated in brackets.

Supplementary Figure 20

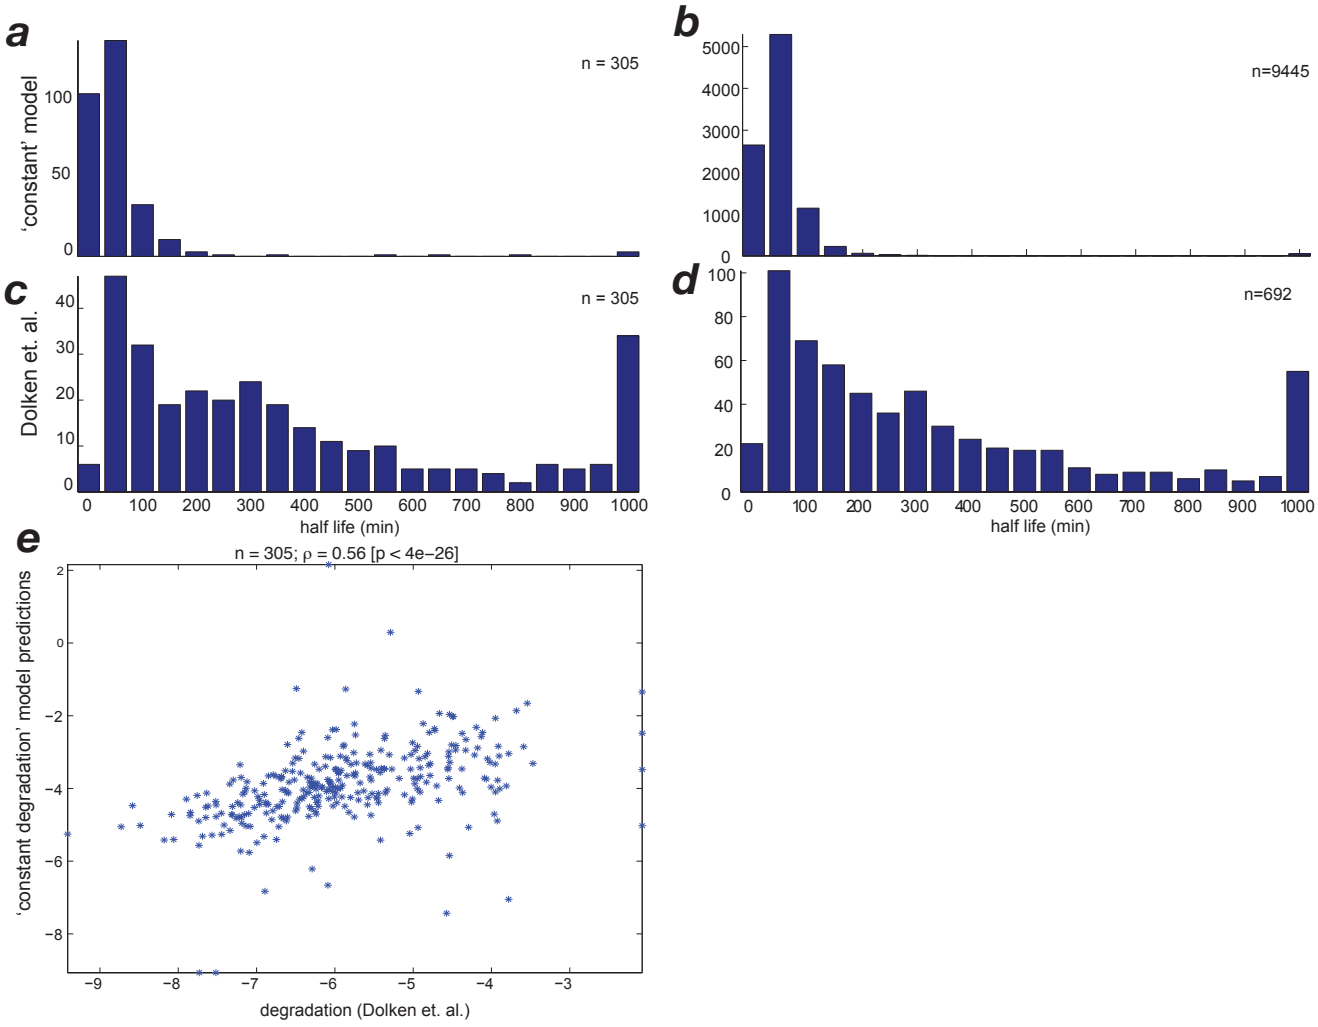

Supplementary Figure 20: Genome wide estimates of (fixed) degradation rates.

Shown are comparisons between RNA degradation rates predicted by the 'constant degradation' model and previously determined rates (Dolken et. al., RNA, 2008) for fibroblasts induced for 1h with interferon, with 1h 4sU labeling and microarray quantification. **(a-d)** Distribution of predicted RNA half-lives using the 'constant degradation' model for 305 genes that appear in both studies (a) and for the complete set of predictions (b), or the published study for 305 genes that appear in both studies (c) and for the complete set of predictions (d). **(e)** Correlation between RNA degradation rates estimates based on the 'constant degradation' model (Y axis, log scale) and RNA degradation rates estimates in the published study (X axis, log scale). Number of genes (n) and Spearman correlation coefficient (ρ) are indicated on top.

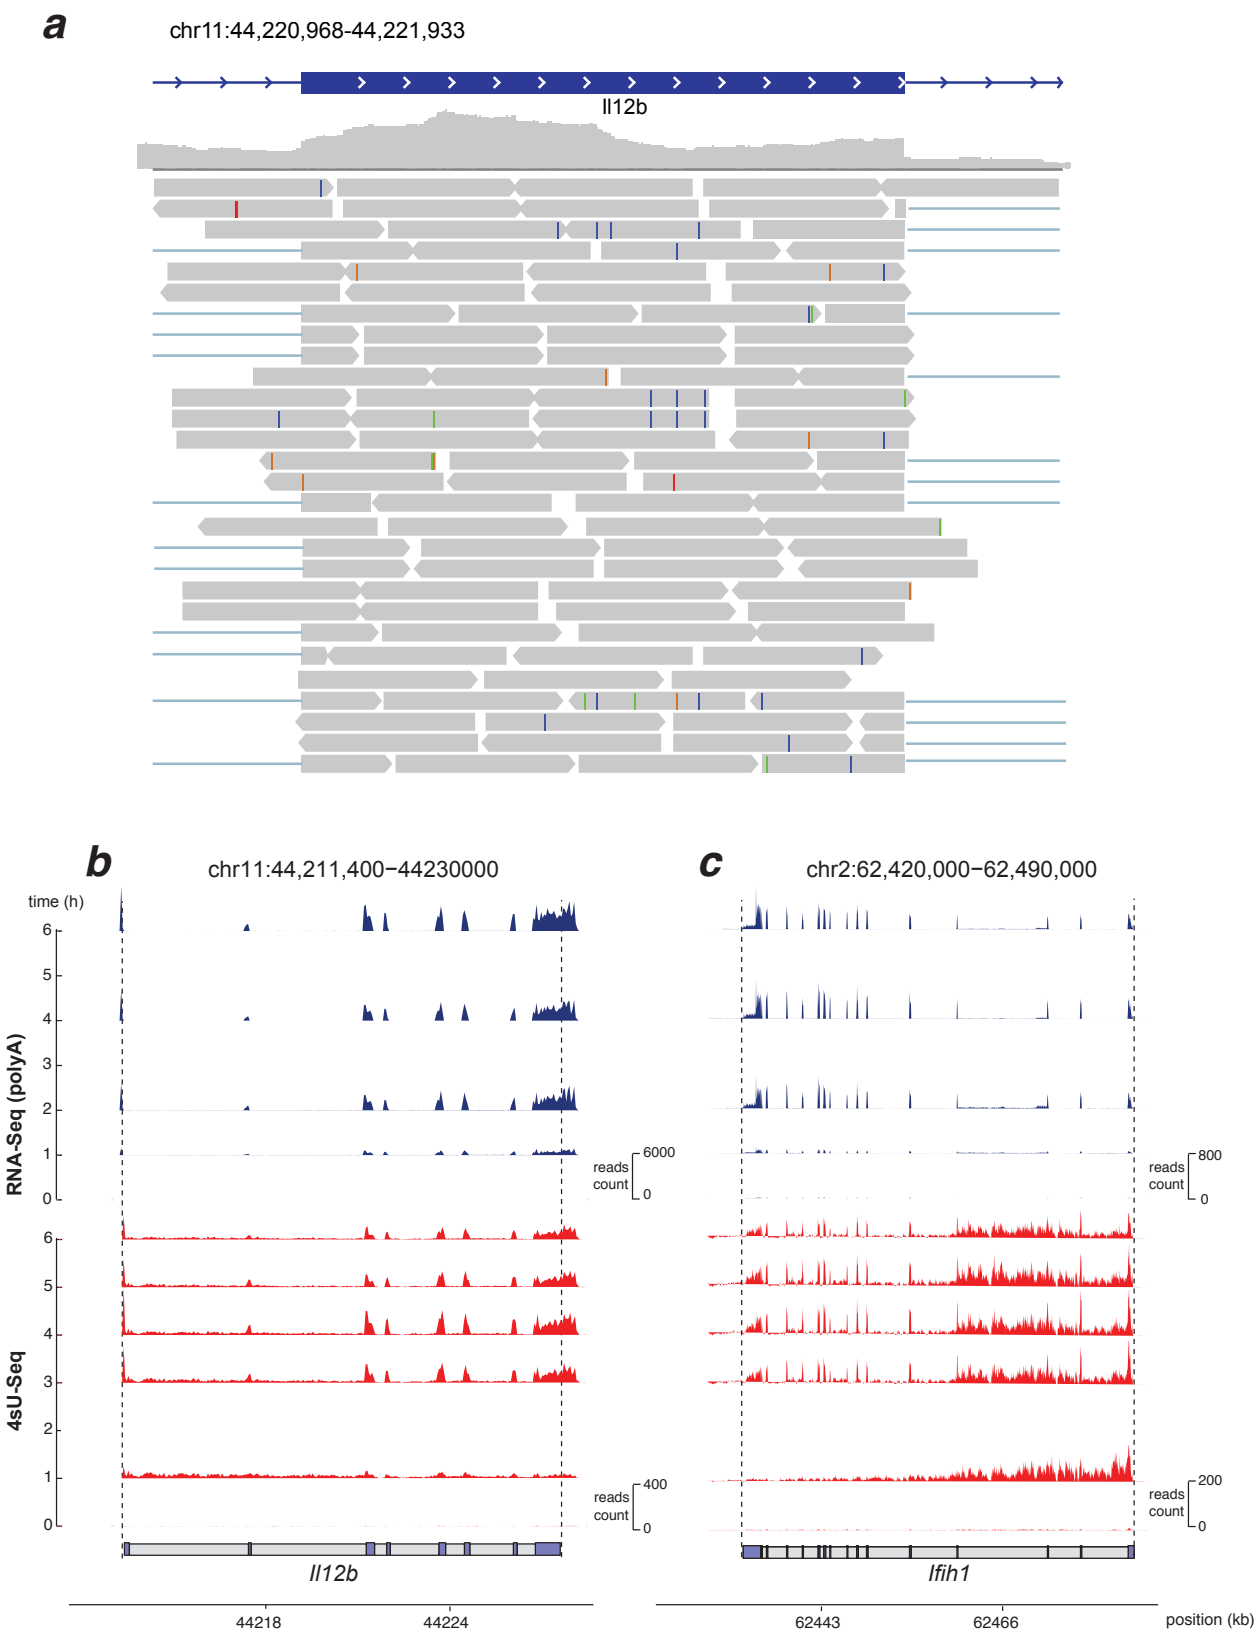
a
chr11:44,220,968-44,221,933
Il12b
b
chr11:44,211,400−44230000
time (h)
RNA-Seq (polyA)
4sU-Seq
reads count
6000
0
400
0
Il12b
44218
44224
c
chr2:62,420,000−62,490,000
800
200
Ifih1
62443
62466
position (kb)

Supplementary Figure 21: Examples of processing rates dynamics.

(a) RNA-4sU-Seq reads map to introns, exons, intron-exon, and exon-exon junctions. Shown is a subset of the reads mapped to the fourth exon of the Il12b gene. Blue track (top): exon boundaries (bold); gray rectangles (bottom tracks): individual reads mapped at the specific position; and colored lines: mismatches in the alignment (blue=C, red=T, yellow=G, green=A). Short rectangles connected to a light blue line indicate that the other end of the read is mapped to an adjacent exon (gapped alignment). **(b-c)** Intronic transcripts are captured by 4sU-Seq but not by RNA-Seq. Shown is the read coverage (height) at each of two example genes (Il12b, b; Ifih1, c) at each time point (track) in each library (bottom: 4sU-Seq, red; top: RNA-Seq, blue). Blue bar: gene locus; light blue boxes: exons.

Supplementary Figure 22

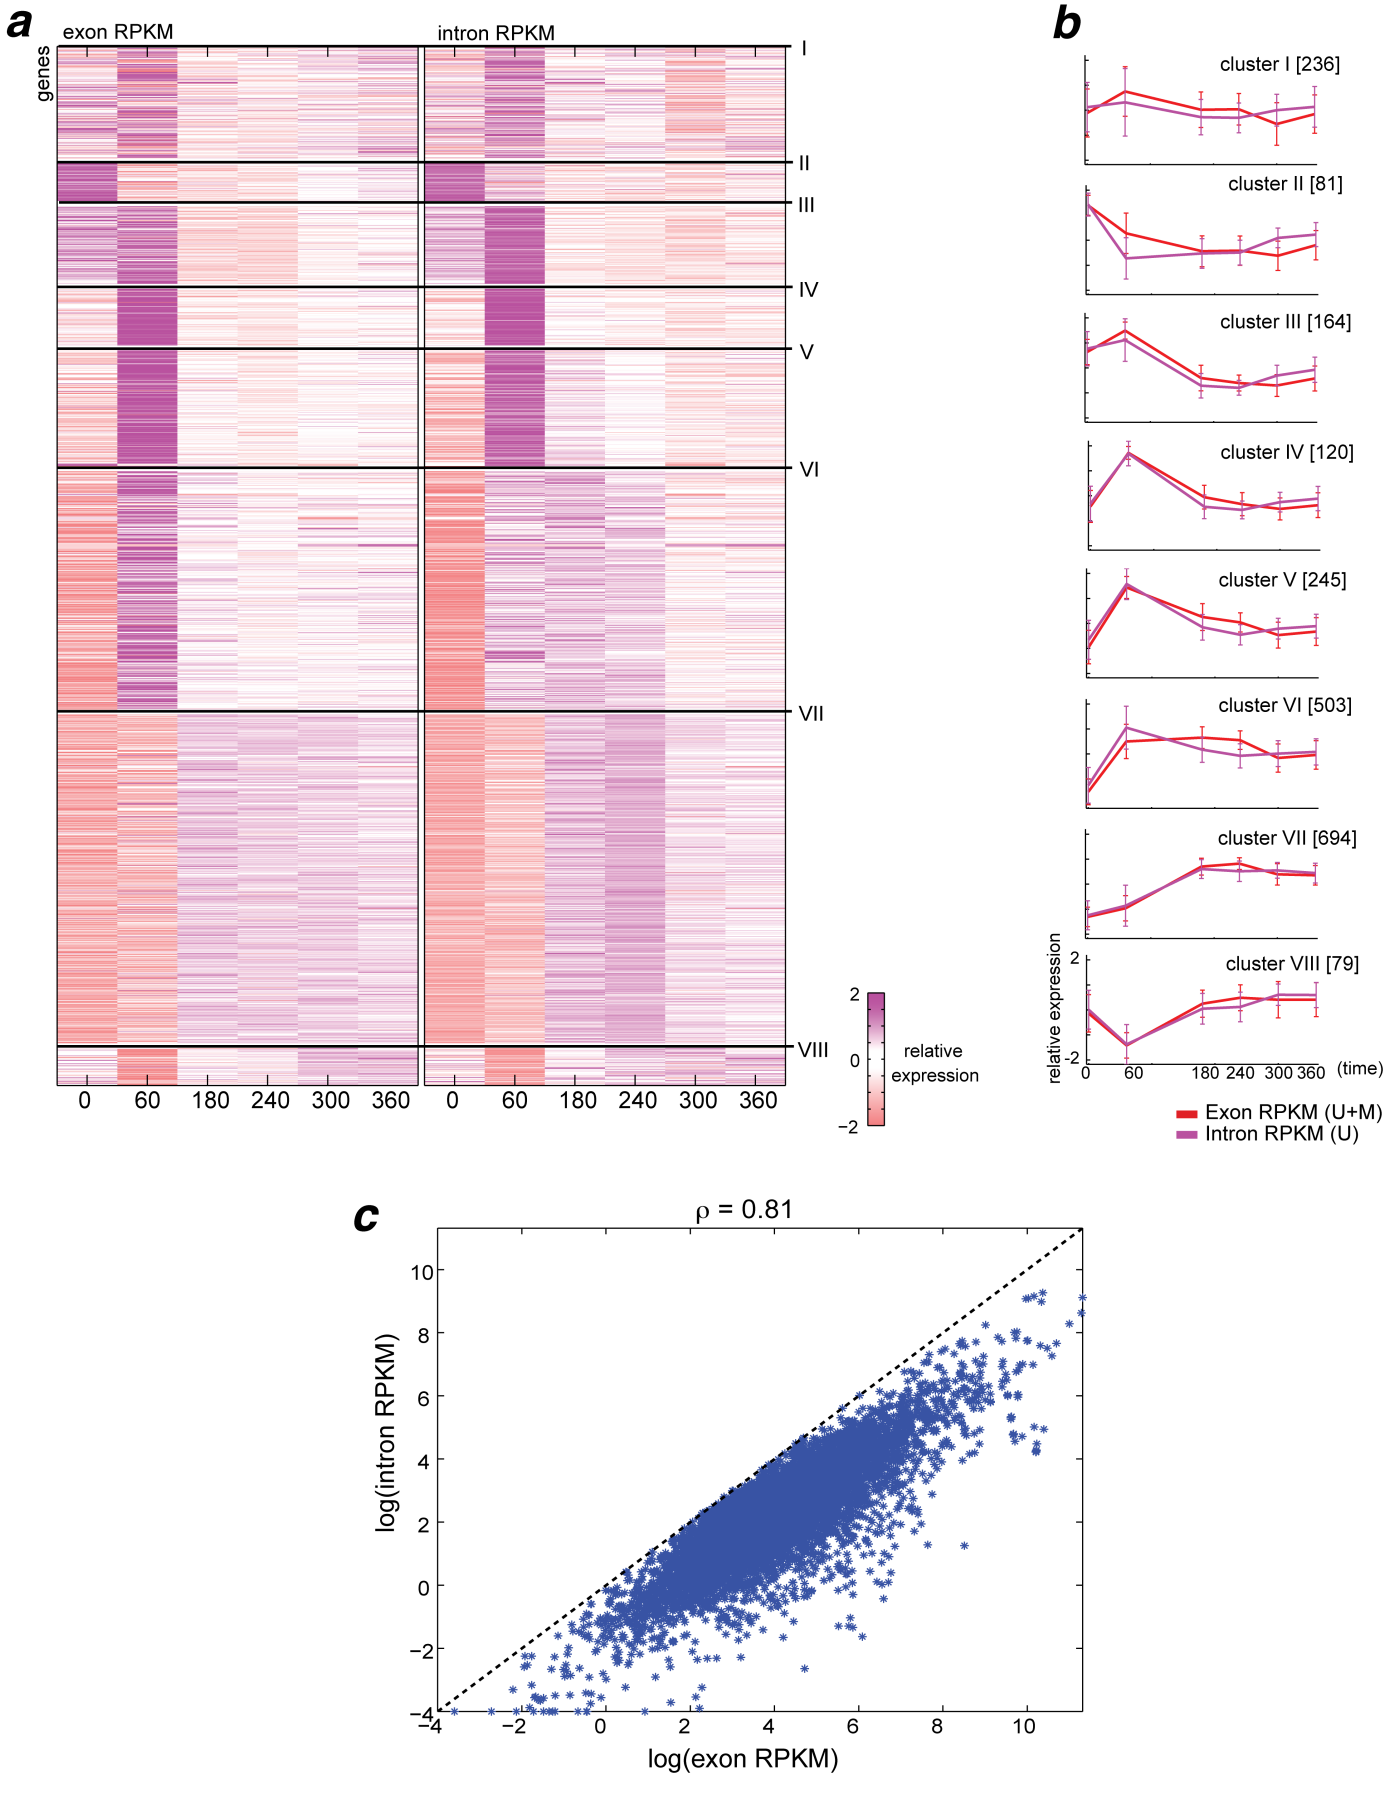

Supplementary Figure 22: Temporal patterns of intronic expression match that of exonic expression across all 8 clusters

(a) Shown are expression profiles (RPKM) of both exons (representing both mature (M) and un-processed (U) transcripts; left) and introns (representing only un-processed (U) transcripts; right) for the 2,122 expressed genes (rows) across 6 time points (columns). Genes are partitioned to 8 expression clusters (I-VIII, right, assigned as in **Supplementary Fig. 16**). Purple: high relative expression; white: mean expression; pink: low relative expression. **(b)** Average profiles for exonic RNA-4sU (red) and intronic RNA-4sU (purple) for each cluster at each time point (0-6h) based on the sequencing data. The size of each cluster is indicated in brackets. **(c)** A high correlation between exon and intron RPKM values. Shown is the correlation between exonic (X axis, log scale) and intronic (Y axis, log scale) RPKM across all 2,122 genes and 6 time points. The Pearson correlation coefficient (ρ) is indicated on top.

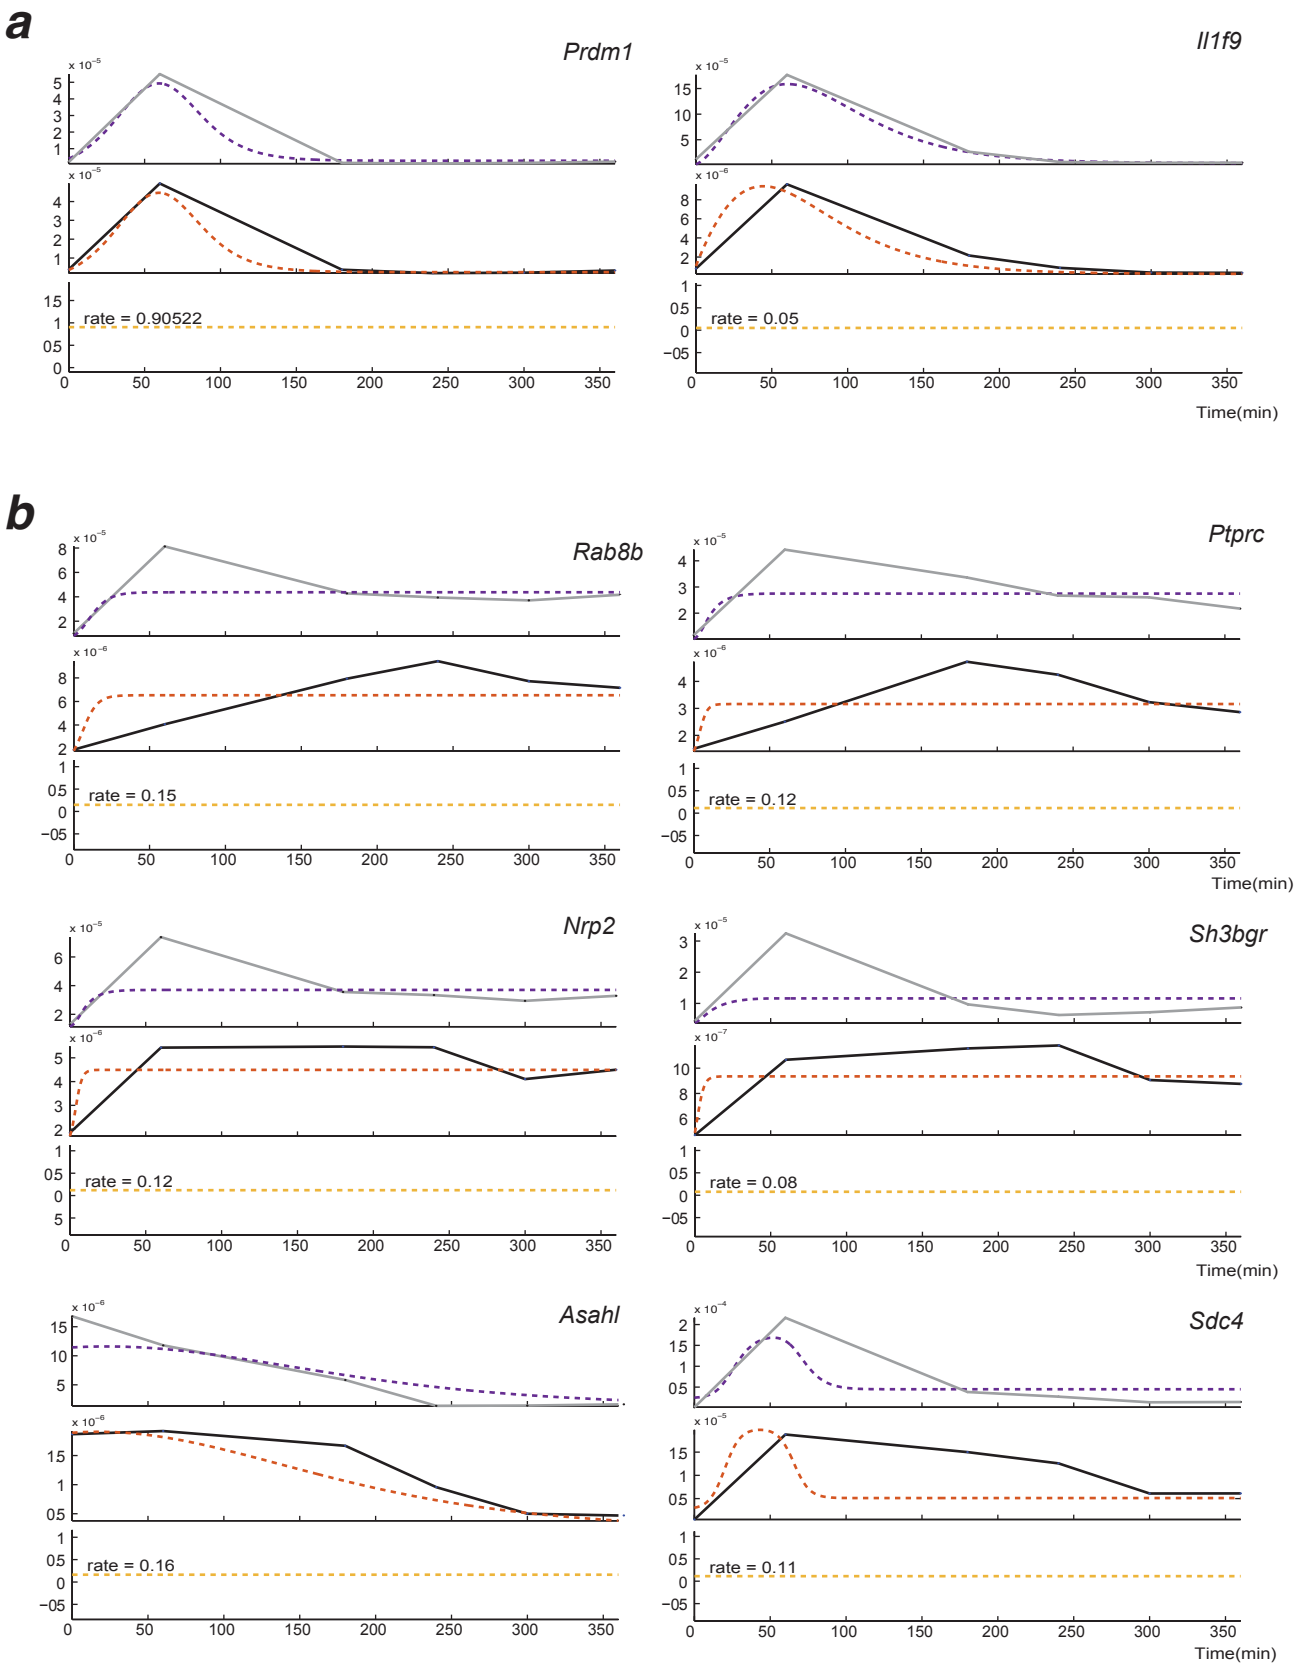
a
Prdm1
rate = 0.90522
Il1f9
rate = 0.05
Time(min)
b
Rab8b
rate = 0.15
Ptprc
rate = 0.12
Time(min)
Nrp2
rate = 0.12
Sh3bgr
rate = 0.08
Time(min)
Asahl
rate = 0.16
Sdc4
rate = 0.11
Time(min)
Measured data:
U (RNA/cell)
U+P (RNA/cell)
Model:
U (RNA/cell)
α(RNA/min*cell)
γ (1/min)

Supplementary Figure 23: The 'constant degradation and processing' model predictions.

Rejection of the 'constant processing and degradation' model might be a result of temporally regulated processing rates. Top: predicted pre-mRNA expression levels (dashed, purple) and measured (solid, gray); middle: predicted transcription rate (dashed, red) and measured exonic RNA-4sU levels (solid, black); bottom: predicted processing rate (solid, orange; value also indicated on top of the line). Gene symbol is indicated on the right. **(a)** Examples of 2 genes that reject the 'constant degradation and processing' model, but show similar expression patterns of pre-mRNA and overall RNA-4sU levels. **(b)** Examples of 6 genes that reject the 'constant processing and degradation' model. Each line shows two examples with a similar behavior, in which pre-mRNA and RNA-4sU (pre-mRNA and mRNA) profiles are different over time.

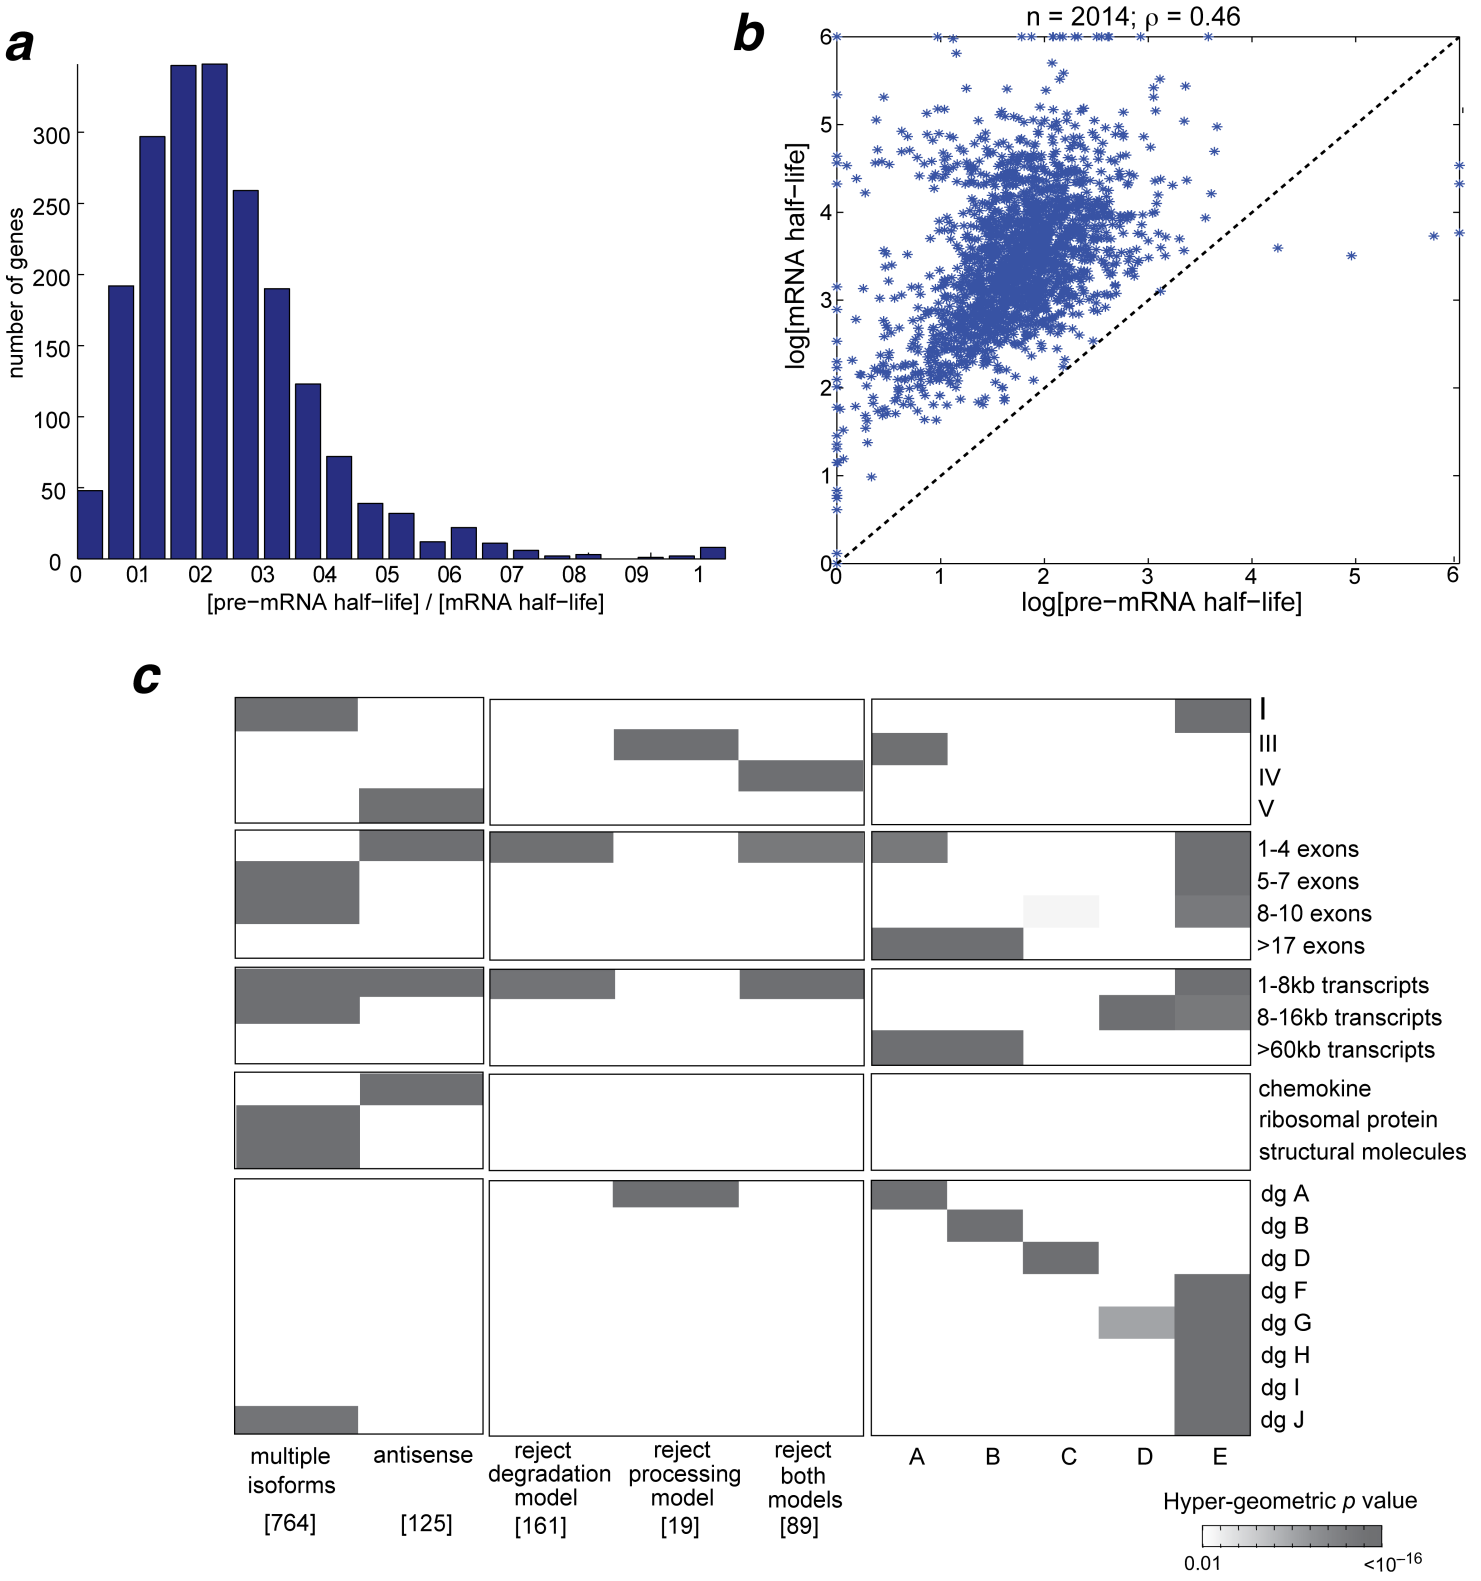
a
number of genes
[pre−mRNA half−life] / [mRNA half−life]
b
n = 2014; ρ = 0.46
log[mRNA half−life]
log[pre−mRNA half−life]
c
I
III
IV
V
1-4 exons
5-7 exons
8-10 exons
>17 exons
1-8kb transcripts
8-16kb transcripts
>60kb transcripts
chemokine
ribosomal protein
structural molecules
dg A
dg B
dg D
dg F
dg G
dg H
dg I
dg J
multiple isoforms [764]
antisense [125]
reject degradation model [161]
reject processing model [19]
reject both models [89]
A
B
C
D
E
Hyper-geometric p value
0.01
<10−16

Supplementary Figure 24: Estimated (constant) processing rates using the 'constant processing and degradation' model.

(a) pre-mRNA half-life is shorter than mRNA half-life. Shown is the distribution of the ratio between pre-mRNA half-life (estimated by the 'constant processing and degradation' model) and the mRNA half-life (estimated by the 'constant degradation' model). **(b)** Significant correlation between mRNA half-life (degradation rate) and pre-mRNA half-life (processing rate). Shown is a scatter plot comparing the RNA processing rates predicted based on the 'constant processing and degradation' model (X axis, log scale) and the RNA degradation rates predicted based on the 'constant degradation' model (Y axis, log scale). The Pearson correlation coefficient is indicated on top. **(c)** Transcripts with multiple isoforms, and transcripts with predicted low or high pre-mRNA half-lives are enriched in functional categories, clusters, exon structures or transcript lengths. Shown are the enrichments (*P*-value, hypergeometric test, grey), of the overlap between the genes with predicted multiple isoforms and anti-sense transcripts (see **Supplementary Notes**, section 6**;** columns); the genes which reject the 'constant processing and degradation' or the 'constant degradation' or both models (columns); and the genes in each of the half-life bins in (**Fig. 4c**) (A-E, columns) and each tested category (rows). Only functional categories with at least one significant enrichment are shown.